

On Threshold Models over Finite Networks*

Elie M. Adam, Munther A. Dahleh, Asuman Ozdaglar[†]

Abstract

We study a model for cascade effects over finite networks based on a deterministic binary linear threshold model. Our starting point is a networked coordination game where each agent's payoff is the sum of the payoffs coming from pairwise interactions with each of the neighbors. We first establish that the best response dynamics in this networked game is equivalent to the linear threshold dynamics with heterogeneous thresholds over the agents. While the previous literature has studied such linear threshold models under the assumption that each agent may change actions at most once, a study of best response dynamics in such networked games necessitates an analysis that allows for multiple switches in actions. In this paper, we develop such an analysis and construct a combinatorial framework to understand the behavior of the model. To this end, we establish that the agents behavior cycles among different actions in the limit and provide three sets of results.

We first characterize the limiting behavioral properties of the dynamics. We determine the length of the limit cycles and reveal bounds on the time steps required to reach such cycles for different network structures. We then study the complexity of decision/counting problems that arise within the context. Specifically, we consider the tractability of counting the number of limit cycles and fixed-points, and deciding the reachability of action profiles. We finally propose a measure of network resilience that captures the nature of the involved dynamics. We prove bounds and investigate the resilience of different network structures under this measure.

*This research is partially supported by an MIT Jacobs Presidential Fellowship, a Siebel Scholarship, a Xerox Fellowship, AFOSR grant FA9550-09-1-0420, and ARO grant W911NF-09-1-0556.

[†]All authors are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139. (emails: eadam@mit.edu, dahleh@mit.edu, asuman@mit.edu)

1 Introduction

Networks intertwine with every aspect of our modern lives, be it through sharing ideas, communicating information, shaping opinions, performing transactions or delivering utilities. Explicitly, we may cite social networks, financial networks, economic networks, communication networks and power networks. Interactions over those many different types of networks require agents to coordinate with their neighbors. In economic networks, technologies that conform to the standards used by other related firms are more productive; in social networks, conformity to the behavior of friends is valuable for a variety of reasons. The desire for such coordination can lead to cascading effects: the adoption decision of some agents can spread to their neighbors and from there to the rest of the network. One of the most commonly used models of such cascading behavior is the linear threshold model introduced by Granovetter [1] to explain a variety of aggregate level behaviors including diffusion of innovation, voting, spread of rumors, riots and strikes. The linear threshold model is also adopted as a model of cascading failures, propagation of diseases and opinion dynamics.

Most analyses of this model in the literature assume that one of the states adopted by the agents (represented by the nodes of a graph) is irreversible, meaning that agents can only make a single switch into this state and can never switch out from it. However, incurring this progressive property in behavior dilutes several perspectives of the dynamics: whereas some situations are best captured by such a variant, many others cannot be captured but by allowing players to revert back to previous actions. One motivation for example would be opinion dynamics in social network: in most situations a player changes opinions back and forth. Another motivation would be recovery within the context of cascading failures or infectious diseases. This said, the literature lacks a satisfactory characterization of the limiting properties of such a model.

In this paper, we consider a model of cascade effects based on binary linear threshold dynamics over finite graphs. We start from an explicit coordination game set over a finite undirected network. The payoff of each agent is the sum of the payoffs from *two player and two action coordination games* the agent plays pairwise with each of the neighbors (the action is fixed across all interactions). We then study the behavior induced by best response dynamics, whereby each agent changes the played action to that which yields highest payoff given the actions of the neighbors. We first establish that best response dynamics are identical to the dynamics traced by the linear threshold model with heterogeneous thresholds for the agents. However, crucially, actions can change multiple times. Thus, the dynamics of interest for the set of problems posed here cannot be studied using existing results and in fact have a different mathematical structure. The main contribution of this paper is to fully characterize these dynamics. We provide three sets of results.

We establish that agent behavior cycles among different actions in the (time) limit, we term such limiting behavior as *convergence cycles*. We characterize the length of those convergence cycles. Ultimately, we show that for any graph structure on the players, any threshold distribution over the players and any initial action configuration played by the players, the limiting behavior of the dynamics get absorbed into action configuration cycles of length at most two. In other words, at the limit, every agent either plays one action, never deviating, or keeps on switching actions at every time step. We then characterize the time required to reach convergence cycles, termed as *convergence time*. Building on the

framework set up, we show that there exists some positive integer c , such that given any graph structure on the players, any threshold distribution over the players and any initial action configuration, the dynamics reach a non-degenerate cycle or a fixed-point in at most cn^2 time steps where n is the number of players. We mention that the work [13] (in the literature on Cellular Automata) considers the same dynamics and proves similar bounds on the length of convergence cycles and a quadratic bound on convergence time. Nevertheless, our proof approach substantially differs in that we emphasize the combinatorial structure of the problem. Our key contribution in this respect lies in two transformations (later defined and termed as *bipartite-expansion* and *symmetric-expansion*) whereby the dynamics (over any initial structure) are extended over a bipartite graph structure where each player has an odd number of neighbors and a majority decision rule. The bipartite-expansion would allow us to decouple our parallel decision scheme to a sequential process, and the symmetric expansion acts as to reduce the proof (of convergence cycle bounds) to counting specific edges in the graph. Using this approach, the quadratic bound on convergence time naturally follows. These transformations would also allow us to answer further questions about fixed points and cycles, and improve the convergence time bound from quadratic to be uniformly not more than the size of the network whenever the graph in concern is either an even-length cycle graph or a tree.

We then study the complexity of counting and decision problems that arise in this model. We are interested in characterizing the number of limiting states the system could get absorbed in. We begin by arguing that no ‘insightful’ uniform upper-bound or lower-bound can be established. Considering only the case of a cycle graph, the number of fixed-points may vary at least from 2 to $2^{n/3}$ depending on the threshold distribution. Instead, we turn to study how tractable it is to count the convergence cycles. We proceed to show that given a graph structure on the player and a threshold distribution over the players as input, the problems of counting the number of limiting configuration classes (i.e. either fixed-points or non-degenerate cycles), counting the number of fixed-points and counting the number of non-degenerate cycles are all #P-Complete. We further show that the problem of deciding whether an action configuration over the network is reachable along the dynamics is NP-Complete and that counting the number of action configuration preceding a *reachable* action configuration is #P-Complete.

Finally, the resilience of networks to invasion by certain types of behavior (e.g., cascades of failures or spread of epidemics) is of central importance in the study of cascades over networks. For the new dynamics defined by our problem, we define a measure of resilience of a network to such invasion that captures the minimal ‘cost of recovery’ needed when the model is confronted with a perturbation in the agents’ action profile. We prove achievable uniform lower-bounds and upper-bounds on the resilience measure, we compute the resilience measure of some network structures and provide basic insight on how different network structures affect this measure.

The paper is related to a large literature on network dynamics and linear threshold models (see e.g., [2]-[8]). A number of papers in this literature investigate the question of whether a behavior initially adopted by a subset of agents (i.e., *the seed set*) will spread to a *large* portion of the network, focusing on the dynamics where agents can make a single switch to one of the behaviors. Morris [2], while starting from a multi-switch version of the dynamics, studied without loss of generality the single-switch version to answer whether there exists a

finite set of initial adopters (in an infinite network with homogeneous thresholds) such that the behavior diffuses to the entire network. In [6], Watts derives conditions for the behavior to spread to a positive fraction of the network (represented by a random graph with given degree distribution) using a branching process analysis. Similarly, Lelarge [7] provides an explicit characterization of the expected fraction of the agents that adopt the behavior in the limit over such networks.

Some works study the single-switch version of the linear threshold model (where nodes never switch back from one of the states) over deterministic finite graphs. Given an initial seed set of adopter, the work [4] characterizes the final set of adopters in terms of *cohesive sets* where cohesion in social groups is measured by comparing the relative frequency of ties among group members to ties with non-members. The work in [5] studies (e.g. in the context of viral marketing) how to target a fixed number of agents (and change their behavior) in order to maximize the spread of the behavior in the network in the (time) limit. It studies the (optimization) problem of maximizing the final set of adopter, under the constraint of picking K initial adopters. It considers various models of cascade, shows that the optimization problem is NP-hard for the linear threshold model, and then provides an algorithm to find an approximation for the optimal set that achieves maximum influence. Although simulating the multi-switch version of the linear threshold model by a single-switch version is both feasible and insightful (see [3] and [5]), the results available (on the single-switch version) do not enable us to characterize the limiting behavior of the richer dynamics.

In the context of network resilience, the recent paper [8] adopts single-switch linear threshold dynamics as a model of failures in a network. This work defines a measure of network resilience that is a function of the graph topology and the distribution over thresholds and studies this measure for different network structures focusing on d -regular graphs (hence ignoring the effect of the degree distribution of a graph on cascaded failures). Here we provide a resilience measure that highlights the impact of heterogeneity in thresholds and degrees of different agents.

Finally, noisy versions of best-response dynamics in networked coordination games were studied in [9] and [10] (see also [11] and [12] in the statistical mechanics literature). The random dynamics in these models can be represented in terms of Markov chains with absorbing states, and therefore do not exhibit the cyclic behavior predicted by the multi-switch linear threshold model studied in this paper.

The rest of the paper proceeds as follows. We begin by a description of the model in Section 2. We then proceed in Section 3 to describe the general behavioral rules of the dynamics. We branch out to characterize convergence cycles and convergence time in Sections 4 and 5. We study the complexity of counting and decision problems in Section 6. We finally propose the network resilience measure in Section 7, and conclude the paper in Section 8. Moreover, certain (non-contiguous) subsets of sections form coherent logical entities. The paper allows readers to select specific topics without having to read through all the sections to ensure understanding. Convergence cycle lengths will be fully treated by the end of Section 4. Readers interested only in convergence time may read Sections 2, 3, 4 (up to 4.3, inclusive) and 5. Readers interested only in the complexity results may restrict their reading to Sections 2, 3, 4 (up to 4.3, inclusive) and 6. Finally, readers interested only in the resilience part may restrict their reading to Sections 2, 3 and 7.

2 Model

The model section consists of two parts. The first sets up a networked coordination game, and characterizes the best-response dynamics in this game. The second proposes a natural extension to the dynamics: it imposes non-equal weights on the pairwise interactions among the players, and allows players to give weights to their own actions.

2.1 The Primary Model

We define a networked coordination game. For a positive integer n , we denote by \mathcal{I}_n the set of n players¹. We define \mathcal{G}_n to be the class of all connected undirected graphs $G(\mathcal{I}_n, E)$ defined over the vertex set \mathcal{I}_n , with edge set E .² To be proper, E is a relation³ on \mathcal{I}_n , but for convenience (since the graph is undirected) we will consider the set E to have cardinality exactly equal to the number of undirected edges. We denote an undirected edge in E by $\{i, j\}$, and we abbreviate it to ij when no confusion arises. For $G(\mathcal{I}_n, E)$ in \mathcal{G}_n , we use $\mathcal{N}_G(i)$ to denote the neighborhood of player i in G , i.e. $\mathcal{N}_G(i) = \{j \in \mathcal{I}_n : ij \in E\}$. We denote by $d_G(i)$ the degree of player i in G , i.e. the cardinality of $\mathcal{N}_G(i)$. We refer to $\mathcal{N}_G(i)$ and $d_G(i)$ respectively as \mathcal{N}_i and d_i when the underlying graph is clear from the context. We finally define \mathcal{Q}_n to be the space of type distributions over the agents, specifically the set of maps from \mathcal{I}_n into $[0, 1]$. We refer to q_i as the type of player i .

Let $\{\mathbb{B}, \mathbb{W}\}$ be a (binary) set of actions, where the symbols \mathbb{B} and \mathbb{W} may be identified with the colors *black* and *white*, respectively. Given a graph $G(\mathcal{I}_n, E)$ in \mathcal{G}_n and a type distribution q in \mathcal{Q}_n , each player i in \mathcal{I}_n plays one action a_i in $\{\mathbb{B}, \mathbb{W}\}$. For $ij \in E$, we define the payoff received by agent i when playing a_i against agent j playing a_j to be

$$g_{i,j}(a_i, a_j) = \begin{cases} q_i & \text{if } a_i = a_j = \mathbb{W} \\ 1 - q_i & \text{if } a_i = a_j = \mathbb{B} \\ 0 & \text{if } a_i \neq a_j \end{cases}. \quad (1)$$

The utility function of player i is the sum of the payoffs from the pairwise interactions with the players in \mathcal{N}_i . Formally, when player j plays action a_j , the utility function of player i is given by:

$$u_i(a_i, a_{-i}) = \sum_{j \in \mathcal{N}_i} g_{i,j}(a_i, a_j), \quad (2)$$

where a_{-i} denotes the action profile of all players except i .

We define \mathcal{A}_n be the space of action profiles⁴ played by the agents, specifically the set of maps from \mathcal{I}_n into $\{\mathbb{B}, \mathbb{W}\}$. The players are assigned an initial action profile \underline{a} , we refer to \underline{a}

¹We use the words player, agent, node and vertex interchangeably. We use the letters i and j to denote agents. We reserve the letter n for the number of players in the game. If it is clear from the context to which set X an element x belongs to, we refrain from mentioning the set X explicitly to simplify notation. Moreover, for any function f with domain \mathcal{I}_n , we will denote $f(i)$ by f_i . In particular, for functions q , k and a with domain \mathcal{I}_n , $q(i)$, $k(i)$ and $a(i)$ are denoted q_i , k_i and a_i respectively.

²For a graph G , we denote by $V(G)$ and $E(G)$ the vertex set and edge set respectively.

³A (binary) relation R on a set A is a subset of $A \times A$. We use the notation aRb to denote $(a, b) \in R$.

⁴We use the words action, assignment and color interchangeably, and use the words profile and configuration interchangeably.

as the action profile of the players at time step 0. For T in \mathbb{Z}^+ ,⁵ every player best responds to the action profile of the players at time step $T - 1$, by choosing the action that maximizes his utility function. We suppose that players play action \mathbb{W} as a tie breaking rule. Formally we impose a strict order on $\{\mathbb{W}, \mathbb{B}\}$ such that $\min\{\mathbb{W}, \mathbb{B}\} = \mathbb{W}$. This tie breaking rule does not affect the behavior of the dynamics whatsoever. However, it does have a natural effect in the network resilience context and this effect will be discussed in Section 7. Suppose we denote by $a_{i,T}$ the action played by player i at time T , then given an initial action configuration \underline{a} in \mathcal{A}_n , for every player i , we recursively define:

$$\begin{aligned} a_{i,0} &= \underline{a}_i \\ a_{i,T} &= \min_{a_i \in \{\mathbb{W}, \mathbb{B}\}} \operatorname{argmax}_{a_{-i,T-1}} u_i(a_i, a_{-i,T-1}), \quad \text{for } T \in \mathbb{Z}^+. \end{aligned} \quad (3)$$

where the min operator breaks ties. The following proposition provides a rule equivalent to that induced by the recursive definition in (3). This characterization is similar to that in [2].

Proposition 2.1. *Let \underline{a} be the initial action configuration, namely the action profile of the players at time step 0. For every positive integer T , player i plays action \mathbb{B} at time step T if and only if (strictly) more than $q_i d_i$ neighbors of player i played action \mathbb{B} at time step $T - 1$.*

Proof. We substitute u_i in (3) with the expressions in (1) and (2), and get that player i plays action \mathbb{B} at time T if and only if

$$\sum_{j \in \mathcal{N}_i} (1 - q_j) \mathbf{1}_{\{\mathbb{B}\}}(a_{j,T-1}) > \sum_{j \in \mathcal{N}_i} q_j \mathbf{1}_{\{\mathbb{W}\}}(a_{j,T-1}),$$

where $\mathbf{1}_\Gamma(x) = 1$ if and only if $x \in \Gamma$. Equivalently, player i plays action \mathbb{B} at time T if and only if

$$\sum_{j \in \mathcal{N}_i} \mathbf{1}_{\{\mathbb{B}\}}(a_{j,T-1}) > q_i d_i.$$

The left-side term is essentially summing the number of neighbors of player i playing action \mathbb{B} . \square

As a technical clarification, we highlight the fact that every player is capable of switching actions both from \mathbb{W} to \mathbb{B} and \mathbb{B} to \mathbb{W} . This contrasts a variant of the dynamics (extensively studied in the literature) where a player can switch out from only one of the actions.

2.2 The Extension Model

Our primary model in the best-response dynamics is such that every node does not take its own action into account and treats the payoffs from the pairwise interactions with equal weights. This corresponds on the part of player i to an unweighted counting of the number of neighbors playing \mathbb{B} at time $T - 1$ to decide whether to play \mathbb{B} at T or not. Our model can take a more general form by allowing each node to play a coordination game with itself, and by assigning symmetric (with respect to the neighbors) weights on the payoffs. We will

⁵We denote by \mathbb{N} the set of non-negative integers, and by \mathbb{Z}^+ the set of positive integers.

also allow those weights to be negative, and thus lose the monotonicity property induced by coordination to get richer dynamics. Neighbors linked by a negatively weighted edge have an incentive to mismatch their actions. In this respect, the induced networked game ceases to be (in general) a networked coordination game. Let $G(\mathcal{I}_n, E)$ be given. Suppose we assign for every ij in E , a non-zero real weight $w_{ij} = w_{ji}$ and for every player i in \mathcal{I}_n , a real weight w_{ii} . Given a q in \mathcal{Q}_n , we extend the utility player i gets to be the weighted sum of the payoffs from the pairwise interactions with the players in $\mathcal{N}_i \cup \{i\}$, specifically when player j plays action a_j ,

$$u_i(a_i, a_{-i}) = \sum_{j \in \mathcal{N}_i \cup \{i\}} w_{ij} g_{i,j}(a_i, a_j).$$

Again, we denote by $a_{i,T}$ the action played by player i at time T . If we let \underline{a} be the initial action configuration, namely $a_{i,0} = \underline{a}_i$, then for every positive integer T , player i plays action $a_{i,T} = \mathbb{B}$ at time step T if and only if

$$\sum_{j \in \mathcal{N}_i \cup \{i\}} w_{ij} \mathbf{1}_{\{\mathbb{B}\}}(a_{j,T-1}) > \theta_i,$$

where we define $\theta_i = q_i \sum_{j \in \mathcal{N}_i \cup \{i\}} w_{ij}$. The primary model described in the previous subsection is then an instance of this model where $w_{ij} = w_{ji} = 1$ for all edges ij in E and $w_{ii} = 0$ for all players i .

3 Description of the Dynamics

We proceed to provide a coarse description of the involved dynamics. We focus only on the primary model throughout this section; all propositions regarding the primary model may be naturally generalized to the extension model. We first highlight what is essential (in our model) for this section. We consider a *finite* set of players \mathcal{I}_n along with three mathematical objects \mathcal{G}_n , \mathcal{Q}_n and \mathcal{A}_n . An element G of \mathcal{G}_n corresponds to the network structure imposed on the players, an element q of \mathcal{Q}_n refers to the type distribution over the players i.e. a function from \mathcal{I}_n into $[0, 1]$, and an element a of \mathcal{A}_n represents an action profile played by the players i.e. a function from \mathcal{I}_n into $\{\mathbb{W}, \mathbb{B}\}$. The elements G , q and a interact as dictated by Proposition 2.1.

3.1 From Types to Thresholds

In Proposition 2.1, Player i uses $q_i d_i$ as a threshold to decide whether to play \mathbb{B} or \mathbb{W} . The value of $q_i d_i$ is non-necessarily an integer, however we may replace it by an integer without modifying the dynamics. To this end, we substitute the set \mathcal{Q}_n by a set \mathcal{K}_n and then modify the statement of Proposition 2.1. We define \mathcal{K}_n to be the space of threshold distributions over the agents, i.e. the set of maps from \mathcal{I}_n into \mathbb{N} . We refer the k_i as the threshold of player i . We make a particular distinction between the word *type* attributed to \mathcal{Q}_n and the word *threshold* attributed to \mathcal{K}_n . Given a pair (G, k) with $k \in \mathcal{K}_n$, Proposition 2.1 generalizes as follows:

Proposition 3.1. *Let \underline{a} be the initial action configuration, namely the action profile of the players at time step 0. For every positive integer T , player i plays action \mathbb{B} at time step T if and only if at least k_i neighbors of player i played action \mathbb{B} at time step $T - 1$.*

The rule in Proposition 3.1 supersedes the rule in Proposition 2.1. For every q in \mathcal{Q}_n there exists a k in \mathcal{K}_n such that $q_i d_i$ may be substituted with the integer k_i for all i without changing the behavior of the players. Thus, a *model* with types may always be simulated by a *model* with integers. However, it is crucial to note that (*strictly*) *more than* is replaced by *at least*. Therefore, in this setting, the converse is not true i.e. not every *model* with integers may be simulated by a *model* with types. Indeed, Proposition 2.1 implies that playing \mathbb{B} is never a best response for player i if no player in \mathcal{N}_i is playing \mathbb{B} . Nevertheless, a player having a *threshold* equal to 0 will always play \mathbb{B} regardless of the actions played by his neighbors. Although we can restrict the thresholds to being non-zero, we refrain from doing so. Instead, we generalize our model (allowing \mathbb{B} to be a best response for i even if no player in \mathcal{N}_i is playing \mathbb{B}) to provide symmetry between both actions \mathbb{B} and \mathbb{W} . We do this for two reasons. The first is to study the linear threshold model as considered in the literature. The second is a technical reason, mainly to ensure closure of the set $\mathcal{G}_n \times \mathcal{Q}_n$ under certain operations e.g. node removal. Nevertheless, any result for the generalized version of the model is inherited by the initial version trivially by inclusion.

Having made the transition from *types* to *thresholds*, we distinguish the nodes having thresholds at the *boundaries* as follows:

Definition 3.2. *Given a pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$, node i in \mathcal{I}_n is called non-valid with respect to (G, k) (or simply non-valid) if k_i is either equal to 0 or (strictly) greater than d_i . A node is called valid if it is not non-valid.*

A non-valid node is then allowed to play only one of the actions in $\{\mathbb{W}, \mathbb{B}\}$ whenever it is allowed to decide on the action to play. Node i will always choose to play \mathbb{B} if its threshold is 0, and will always choose to play \mathbb{W} if its threshold is (strictly) greater than d_i .

Finally, given a pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$, we denote by G_k the map from \mathcal{A}_n into \mathcal{A}_n such that for player i , $(G_k a)_i = \mathbb{B}$ if and only if at least k_i players are in $a^{-1}(\mathbb{B}) \cap \mathcal{N}_i$.⁶ From this perspective, given an initial configuration \underline{a} in \mathcal{A}_n , the sequence $\underline{a}, G_k \underline{a}, G_k^2 \underline{a}, \dots$ corresponds to the sequence of action profiles $\underline{a}, a_1, a_2, \dots$ where $a_T = G_k^T \underline{a}$ is the action profile played by the players at time T if they act in accordance with the rule in Proposition 3.1.

3.2 The Limiting Behavior

To understand the limiting behavior, we note two fundamental properties: the space \mathcal{A}_n has finite cardinality, and Proposition 3.1 is deterministic. Since \mathcal{A}_n is finite, if we let a_0, a_1, a_2, \dots be any infinite sequence of action profiles played by the agents according to Proposition 3.1, then there exists at least one action profile \hat{a} that will appear infinitely many times along this sequence. Since the dynamics are deterministic (and that actions at time $T+1$ depend only on actions at time T), the same sequence of action profiles appears

⁶Let $f : A \rightarrow B$ and $g : B \rightarrow C$ be functions, we denote by gf the function $g \circ f : A \rightarrow C$. In particular, if a function f maps a set A to itself, for a non-negative integer m , we denote by f^m the function $f \circ f^{m-1} \circ \dots \circ f^0$ where f^0 is the identity map on A .

between any two consecutive occurrences of \hat{a} . This means that after a finite time step, the sequence of action profiles will cycle among action profiles.

Let us consider a different representation of the dynamics. Given a pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$, we define a (binary) relation \rightarrow on \mathcal{A}_n such that for a and b in \mathcal{A}_n , $a \rightarrow b$ if and only if $b = G_k a$. The graph $H(\mathcal{A}_n, \rightarrow)$ then forms a directed graph (possibly with self loops) on the vertex set taken to be the space of action profiles \mathcal{A}_n , and an action profile a is connected to an action profile b by a directed edge (a, b) going from a to b if and only if $b = G_k a$. Suppose we pick a vertex a , namely an action configuration, and perform a walk on vertices along the edges in H starting from a . The walk eventually cycles vertices in the same order. Every initial action profile leads to one cycle, and two action profiles need not lead to the same cycle. We formalize the idea in the following definitions.

Definition 3.3. *Given (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$, for two action profiles a and b in \mathcal{A}_n , we say that a can be reached from b with respect to G_k if there exists a non-negative integer T such that $a = G_k^T b$. Formally, we define the relation \mathcal{R}_{G_k} on \mathcal{A}_n such that for a and b in \mathcal{A}_n , $a \mathcal{R}_{G_k} b$ if and only if there exists a non-negative integer T such that $a = G_k^T b$.*

For a and b in \mathcal{A}_n , we have $a \mathcal{R}_{G_k} b$ if and only if there exists a directed path in $H(\mathcal{A}_n, \rightarrow)$ from vertex b to vertex a . The graph H is not necessarily weakly-connected,⁷ and by the argument provided at the beginning of this subsection, every weakly-connected component of H necessarily contains a directed cycle (possibly a self-loop). Moreover, each vertex in H can have at most one outgoing edge, therefore every weakly-connected component of H cannot contain more than one directed cycle.⁸ We formally characterize the set of those cycles as follows:

Definition 3.4. *Given a pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$, we define $CYCLE_n(G, k)$ to be the collection of subsets of \mathcal{A}_n , such that for every C in $CYCLE_n(G, k)$, if a and b are in C then we have both $a \mathcal{R}_{G_k} b$ and $b \mathcal{R}_{G_k} a$, and for every c in $\mathcal{A}_n \setminus C$, there does not exist an action configuration a in C such that $a \mathcal{R}_{G_k} c$. We refer to the elements of $CYCLE_n(G, k)$ as convergence cycles.*

The condition “ $a \mathcal{R}_{G_k} b$ and $b \mathcal{R}_{G_k} a$ ” can be concisely replaced by “ $a \mathcal{R}_{G_k} b$ ”, however we keep it as such to stress on the fact that both a can be reached from b and b can be reached from a . The second condition ensures that C is in $CYCLE_n(G, k)$ only if there exists no larger cycle C' containing C .

Cycles in $CYCLE_n(G, k)$ consisting of only one action configuration are fixed-points of G_k and so will be referred to as *fixed-points*. Cycles in $CYCLE_n(G, k)$ consisting of more than one action configuration will be referred to as *non-degenerate cycles* (as opposed to *fixed-points* which are *degenerate cycles*).

In the remaining part of this section, we provide a broad overview of the main results (without proofs). We only mention results concerning the dynamical properties of the model

⁷A directed graph G is said to be weakly-connected if for any vertices u and v in the vertex set of G , there exists an *undirected* path connecting u to v . A weakly-connected component of G is a maximal subgraph of G that is weakly-connected.

⁸If we construct a relation \mathcal{C} on \mathcal{A}_n such that for a and b in \mathcal{A}_n , $a \mathcal{C} b$ if and only if $a \mathcal{R}_{G_k} b$ or $b \mathcal{R}_{G_k} a$, then \mathcal{C} is an equivalence relation on \mathcal{A}_n . In this setting, two configurations in \mathcal{A}_n are in the same equivalence class with respect to the relation \mathcal{C} if and only if they are in the same weakly-connected component in H .

(e.g. the length of convergence cycles, convergence time, the number of fixed points or cycles and reachability). We do not consider yet any resilience measure or bounds thereof, we do so in Section 7.

3.3 An Overview of Convergence Results

Given the limiting cyclic behavior, the most natural starting point would be to characterize the length of the cycles in the equivalence classes as a function of the imposed graph structure and the threshold distribution.

Theorem 3.5. *For every positive integer n , every (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ and every cycle C in $\text{CYCLE}_n(G, k)$, the cardinality of C is less than or equal to 2.*

Put differently, given a network structure G , a threshold distribution k and an initial action profile a , if we iteratively apply G_k on a ad infinitum to get a sequence of best response action profiles, along the sequence of actions considered by player i , player i will eventually either settle on playing one action, or switch action on every new application of G_k .

We further show (in Sections 4.4 and 4.5) that such a limiting behavior also holds for the extension model. We then proceed to characterize the number of iterations needed to reach a convergence cycle.

Definition 3.6. *For every positive integer n , and every (G, k, a) in $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$, we define $\delta_n(G, k, a)$ to be equal to the smallest non-negative integer T such that there exists a cycle C in $\text{CYCLE}_n(G, k)$ and b in C with $G_k^T a = b$.*

The quantity $\delta_n(G, k, a)$ denotes to the minimal number of iterations needed until a given action configuration a reaches a cycle, when iteratively applying G_k . We refer to $\delta_n(G, k, a)$ as the *convergence time* from a under G_k .

Theorem 3.7. *For some positive integer c , every positive integer n , and every (G, k, a) in $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$, the convergence time $\delta_n(G, k, a)$ is less than or equal to cn^2 .*

We further improve the results on convergence time and get a bound that is linear in the size of the network when the graphs are restricted to cycle graphs or trees.

Theorem 3.8. *For all positive integers n , and every (G, k, a) in $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$ where G is an even-length cycle graph or a tree, the convergence time $\delta_n(G, k, a)$ is less than or equal to n .*

A linear bound can also be derived for complete graphs, but we do not consider this case in this paper. Instead, we refer the reader to [18].

We proceed to characterize the number of fixed-points and (non-degenerate) convergence cycles and present an overview of results on decision and counting problem that arise within this framework.

3.4 An Overview of Complexity Results

We first argue in Section 6 that no insightful *uniform* bound on the number of limiting configurations exists; the range of such a bound will be too large. In this case, how well can we characterize the number of fixed-points and non-degenerate convergence cycles? To this end, we study the complexity of counting those numbers. We consider the counting problems⁹ $\#CYCLE$, $\#FIX$ and $\#2CYCLE$ that take $\langle n, G, k \rangle$ as input, where n is a positive integer and (G, k) belongs to $\mathcal{G}_n \times \mathcal{K}_n$, and outputs the cardinality of $CYCLE_n(G, k)$, the number of fixed points and the number of non-degenerate cycles respectively.

Theorem 3.9. $\#CYCLE$, $\#FIX$ and $\#2CYCLE$ are $\#P$ -Complete.

One has to be subtle towards what such result entails. This result does not imply that no characterization of the number of fixed-points or non-degenerate cycles is possible whatsoever, but rather that we would be unable to get an arbitrarily refined characterization of that number assuming $\#P$ is not in FP.

We further show in Section 6.1 that those counting problems remain hard even if we restrict the graphs to be bipartite and impose homogeneous thresholds on the players.

We then proceed to the question of reachability whereby given a graph structure G , a type distribution k we decide whether a certain action configuration can be reached from some other configuration. We define the language PRED to consist of all 4-tuples $\langle n, G, k, a \rangle$, where n is a positive integer, (G, k, a) belongs to $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$ with $G_k(a)^{-1} \neq \emptyset$.

Theorem 3.10. PRED is NP-Complete.

Given a graph structure G , a type distribution k and a configuration a , computing the number of configurations b from which a can be reached by applying G_k only once on b is then necessarily hard. Instead, we define the counting problem $\#reachable\text{-PRED}$ to take $\langle n, G, k, a \rangle$ as input, where n is a positive integer, (G, k, a) and $G_k(a)^{-1} \neq \emptyset$ and output the cardinality of $G_k^{-1}(a)$. We get the following result:

Theorem 3.11. $\#reachable\text{-PRED}$ is $\#P$ -Complete.

The results are derived from thresholds in \mathcal{K}_n instead of types of \mathcal{Q}_n . However, the results trivially extend to types as follows: Convergence results hold by inclusion; complexity results hold since they still hold if we restrict (G, k) to contain no non-valid node. We devote Sections 4 and 5 to convergence results, and Section 6 to complexity results.

4 On Convergence Cycles

We begin by studying the following problem: given a graph G in \mathcal{G}_n and a threshold distribution k in \mathcal{K}_n , how many action configurations does a cycle in $CYCLE_n(G, k)$ contain? Ultimately, we show that for any graph and any threshold distribution, the cycles in $CYCLE_n(G, k)$ consist of at most two action configurations. We begin the analysis by considering cycle graphs¹⁰, then proceed to trees. We refer the reader to [18] for an analysis on

⁹We refer the reader to Appendix A for a short review of the required background.

¹⁰A cycle graph is a 2-regular connected graph, we use both terms interchangeably.

complete graphs. Each of those special cases is treated by exploiting its graphic properties. Obviously, most of those properties are not shared among all graphs, and some cannot even be generalized to general graphs. Nevertheless, we explicitly provide results over those toy examples to build up the intuition of the reader and construct the combinatorial framework slowly as we go along. After trees, we consider general graphs. We then generalize the results to the extension model.

4.1 Cycle Graphs

Let us consider a pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ where G is a path, i.e. every agent is connected to at most two other agents and no cycles in the graphs are allowed. Our first intent is to characterize the length of the limiting cycles in that case. Suppose k is picked in such a way that some players are non-valid¹¹, then we know that those players can only play one action after some finite time step. With respect to the analysis concerned, we may remove those players, update the thresholds of the neighboring players accordingly and end up with a collection of disconnected paths. Restricting the analysis to one of the paths leads us back to the initial case. Therefore, we will assume that every node in the graph is valid: this implies that k is equal to 1 for the nodes having degree 1 and k takes values in $\{1, 2\}$ for the nodes having degree equal to 2. Moreover, to take care of the boundary case, we will connect the 1-degree nodes together, and so forming a ring of agents. The graph in consideration is then the 2-regular connected graph. We then relax k to take values in $\{1, 2\}$ over \mathcal{I}_n .

Given that the thresholds of the nodes are either 1 or 2, it is useful to explicitly state the decision rules as follows. Let a be some action configuration in \mathcal{A}_n , if node i has a threshold k_i equal to 1, then node i is \mathbb{B} in $G_k a$ if and only if either one of its neighbors is \mathbb{B} in a . Similarly, if node i has a threshold k_i equal to 2, then node i is \mathbb{B} in $G_k a$ if and only if both of its neighbors are \mathbb{B} in a .

Let us impose a strict ordering on $\{\mathbb{W}, \mathbb{B}\}$ such that $\min\{\mathbb{W}, \mathbb{B}\} = \mathbb{W}$. This translates to $\mathbb{W} \wedge \mathbb{B} = \mathbb{W}$ for notational convenience.¹² For elements a, b and c in $\{\mathbb{W}, \mathbb{B}\}$, the following identities can be checked:

$$\begin{array}{ll} a \wedge a = a & a \vee a = a \\ a \wedge \mathbb{B} = a & a \vee \mathbb{W} = a \\ a \wedge b = b \wedge a & a \vee b = b \vee a \\ a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c) & a \vee (b \wedge c) = (a \vee b) \wedge (a \vee c) \\ a \wedge (b \wedge c) = (a \wedge b) \wedge c & a \vee (b \vee c) = (a \vee b) \vee c. \end{array}$$

Given a pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ where G is 2-regular and k takes values in $\{1, 2\}$, we define a map τ from \mathcal{I}_n into $\{\vee, \wedge\}$ such that $\tau_i = \vee$ if and only if $k_i = 1$. Similarly, let us define two maps s and p from \mathcal{I}_n into \mathcal{I}_n (we refer to them as successor and predecessor) such that i and s_i are neighbors, i and p_i are neighbors and $(sp)_i = (ps)_i = i$. The dynamics are then

¹¹See definition 3.2

¹²For x_1 and x_2 in a strictly ordered set, we denote $\max\{x_1, x_2\}$ and $\min\{x_1, x_2\}$ by $x_1 \vee x_2$ and $x_1 \wedge x_2$ respectively.

represented as follows:

$$(G_k a)_i = a_{p_i} \tau_i a_{s_i}.$$

We give a quick example to illustrate. Let us consider a 2-regular connected graph G over the set \mathcal{I}_n for $n \geq 5$ and suppose we are given a threshold distribution k in \mathcal{K}_n taking values in $\{1, 2\}$. Let us choose a node i from \mathcal{I}_n . The nodes s_i and p_i are then (distinct) neighbors of node i . Let a be an action configuration in \mathcal{A}_n and suppose $k_i = 1$, then $(G_k a)_i = \mathbb{B}$ if and only if either $a_{s_i} = \mathbb{B}$ or $a_{p_i} = \mathbb{B}$ i.e. at least one neighbor is \mathbb{B} . We can rewrite the previous statement as:

$$(G_k a)_i = \max\{a_{s_i}, a_{p_i}\} = a_{s_i} \vee a_{p_i} = a_{s_i} \tau_i a_{p_i}.$$

This follows from the fact that we imposed a strict ordering on $\{\mathbb{W}, \mathbb{B}\}$ such that $\min\{\mathbb{W}, \mathbb{B}\} = \mathbb{W}$, and defined $\tau_i = \vee$ if and only if $k_i = 1$. Furthermore, s_i has both i and $(ss)_i$ as neighbors. Suppose $k_{s_i} = 2$, then $(G_k a)_{s_i} = \mathbb{B}$ if and only if both $a_i = \mathbb{B}$ and $a_{(ss)_i} = \mathbb{B}$ i.e. at least two neighbor are \mathbb{B} . Similarly, we can rewrite the previous statement as:

$$(G_k a)_{s_i} = \min\{a_i, a_{(ss)_i}\} = a_i \wedge a_{(ss)_i} = a_i \tau_{s_i} a_{(ss)_i}.$$

To conclude the example, we further suppose that $k_{p_i} = 2$. The node p_i has both i and $(pp)_i$ as neighbors, and similarly to the rule of node s_i , we have $(G_k a)_{p_i} = a_i \wedge a_{(pp)_i}$. We may now express $(G_k^2 a)_i$ in terms of actions in a as follows:

$$(G_k^2 a)_i = (G_k(G_k a))_i = (G_k a)_{s_i} \vee (G_k a)_{p_i} = (a_i \wedge a_{(ss)_i}) \vee (a_i \wedge a_{(pp)_i}).$$

Using distributivity (as defined in the identities earlier) we get:

$$(a_i \wedge a_{(ss)_i}) \vee (a_i \wedge a_{(pp)_i}) = a_i \wedge (a_{(ss)_i} \vee a_{(pp)_i}).$$

We now generalize the last part of the example. We consider a pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ where G is a cycle graph and k takes values in $\{1, 2\}$. In this setting, $(G_k^2 a)_i$ would only depend on the actions of nodes $(ss)_i$, $(pp)_i$ and i itself in a . In particular, there are a total of eight possible decision rules, we summarize them in the following table:

	τ_{p_i}	τ_i	τ_{s_i}	$(G_k a)_i$	$(G_k^2 a)_i$
c.1	\vee	\vee	\vee	$a_{p_i} \vee a_{s_i}$	$a_i \vee (a_{(ss)_i} \vee a_{(pp)_i})$
c.2	\vee	\vee	\wedge	$a_{p_i} \vee a_{s_i}$	$a_i \vee a_{(pp)_i}$
c.3	\vee	\wedge	\vee	$a_{p_i} \wedge a_{s_i}$	$a_i \vee (a_{(ss)_i} \wedge a_{(pp)_i})$
c.4	\vee	\wedge	\wedge	$a_{p_i} \wedge a_{s_i}$	$a_i \wedge a_{(ss)_i}$
c.5	\wedge	\vee	\vee	$a_{p_i} \vee a_{s_i}$	$a_i \vee a_{(ss)_i}$
c.6	\wedge	\vee	\wedge	$a_{p_i} \vee a_{s_i}$	$a_i \wedge (a_{(ss)_i} \vee a_{(pp)_i})$
c.7	\wedge	\wedge	\vee	$a_{p_i} \wedge a_{s_i}$	$a_i \wedge a_{(pp)_i}$
c.8	\wedge	\wedge	\wedge	$a_{p_i} \wedge a_{s_i}$	$a_i \wedge (a_{(ss)_i} \wedge a_{(pp)_i})$

We proceed by defining strong assignments, then state a first proposition on the dynamics induced by the above table.

Definition 4.1. For any positive integer n greater than 1, any graph G in \mathcal{G}_n , every threshold distribution k in \mathcal{K}_n , an action c is called a strong action (or strong assignment) for player i in \mathcal{I}_n if once played by player i at time step T , it is played by player i at time step $T + 2m$ for all positive integers m , regardless of what is played by the neighbors of node i .

Given the definition, every node in the setting concerned in this subsection has a strong assignment.

Proposition 4.2. For any integer n greater than 2, any 2-regular graph G in \mathcal{G}_n , any threshold distribution k in \mathcal{K}_n taking values in $\{1, 2\}$, every player i in \mathcal{I}_n has a strong assignment.

Proof. Let i be a player in \mathcal{I}_n , to prove the result it would be enough to investigate the update rule over two iterated applications of G_k of player i , i.e. the value $(G_k^2 a)_i$ takes. Referring back to the previous table, notice that in the case of each threshold distribution over $\{p_i, i, s_i\}$, for $(G_k^2 a)_i$ to be equal to \mathbb{B} , the condition $a_i = \mathbb{B}$ is either sufficient or necessary. In the case where $a_i = \mathbb{B}$ is sufficient, if $a_i = \mathbb{B}$ then $(G_k^{2m} a)_i = \mathbb{B}$ for all positive integers m , and so \mathbb{B} is a strong assignment for player i . Likewise, in the case where $a_i = \mathbb{B}$ is necessary, if $a_i = \mathbb{W}$ then $(G_k^{2m} a)_i = \mathbb{W}$ for all positive integers m , and so \mathbb{W} is a strong assignment for player i . \square

In particular, c.1, c.2, c.3 and c.5 correspond to \mathbb{B} being a strong assignment for player i , and c.4, c.6, c.7 and c.8 correspond to \mathbb{W} being a strong assignment for player i . We now characterize the length of the convergence cycles.

Proposition 4.3. For any integer n greater than 2, any 2-regular graph G in \mathcal{G}_n and any threshold distribution k in \mathcal{K}_n taking values in $\{1, 2\}$, each cycle C in $CYCLE_n(G, k)$ has cardinality less than or equal to 2.

Proof. Let a be an action configuration in \mathcal{A}_n and suppose we construct the sequence $a, G_k a, G_k^2 a, G_k^3 a, \dots$. For notational convenience, let us denote $G_k^T a$ by a^T . We consider the subsequence a^0, a^2, a^4, \dots , choose a player i in \mathcal{I}_n and then observe the evolution of the action played by player i over two time steps, i.e. we consider the sequence $a_i^0, a_i^2, a_i^4, \dots$. Without any loss of generality, let us assume that \mathbb{B} is the strong assignment. Either \mathbb{B} appears in the sequence or \mathbb{B} does not appear in the sequence. If it does appear, then there exists a positive integer M such that $a_i^{2m} = \mathbb{B}$ for all $m \geq M$. If it does not appear, then $a_i^{2m} = \mathbb{W}$ for all non-negative integers m . Either way, for every player i , there exists a non-negative integer T_i and an action c in $\{\mathbb{W}, \mathbb{B}\}$ such that $a_i^{2m} = c$ for all $m \geq T_i$. If we set $T = \max_i T_i$, then there exists an action profile \hat{a} such that $a^{2m} = \hat{a}$ for all $m \geq T$. We then get that $\{\hat{a}, G_k \hat{a}\}$ is the cycle reached from a . It follows that if we let C be any cycle in $CYCLE_n(G, k)$ and we let a be an action configuration in C , then necessarily $C = \{a, G_k a\}$. \square

We transition to investigate the behavior when the graph structure is a tree.

4.2 Trees

We consider in this section dynamics on trees, namely acyclic connected graphs. In this section, the letter T shall always be used to denote trees, and never time as was done sometimes in previous sections. Given a tree T in \mathcal{G}_n , if we label a node r in \mathcal{I}_n as *root*, the children of node i (with respect to the root r) are all the neighbors of i , not lying on the path from the root r to node i . Finally, a leaf in the tree T is a node having degree 1. Fortunately, strong assignments appear in the dynamics on trees. We begin by stating the following proposition:

Proposition 4.4. *For any integer n greater than 1, any tree T in \mathcal{G}_n and any threshold distribution k in \mathcal{K}_n such that all nodes are valid with respect to (G, k) , pick a root r for the tree, then for every node i where all its children (with respect to r) are leaves, i has at least one strong assignment. In particular, if $k_i \geq 1$, then \mathbb{W} is a strong assignment and if $k_i < d_i$ then \mathbb{B} is a strong assignment.*

Proof. We know that for each node i where all its children (with respect to r) are leaves, node i has at least $d_i - 1$ leaves connected to it. Since all nodes are considered to be valid, then each leaf has a threshold of 1. Suppose $k_i > 1$, it then follows that $(G_k^2 a)_i = a_i \wedge \phi(a)$ for some map ϕ from \mathcal{A}_n into $\{\mathbb{W}, \mathbb{B}\}$. Suppose $k_i < d_i$, it then follows that $(G_k^2 a)_i = a_i \vee \phi(a)$ for some map ϕ from \mathcal{A}_n into $\{\mathbb{W}, \mathbb{B}\}$. \square

In this case, note that if $1 < k_i < d_i$, then i has both \mathbb{B} and \mathbb{W} as strong assignment. This fact implies that i will never change its color over two time steps. We note that the proposition considers only the case where all nodes in concern are valid.

Aside being acyclic, trees enjoy bipartiteness: a crucial property that will be heavily relied on when considering general graphs. We begin to convey how the bipartite property of graphs may be exploited. The definitions and results to follow apply to general bipartite graphs.

Definition 4.5. *Let P be a subset of \mathcal{I}_n , for (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$, we define $G_k|_P$ to be the restriction of G_k to act on the actions of the players in P . Formally, for a in \mathcal{A}_n ,*

$$(G_k|_P a)_i = \begin{cases} (G_k a)_i & \text{if } i \in P \\ a_i & \text{if } i \notin P \end{cases}$$

We note that we are not restricting the domain of the function, $G_k|_P$ is indeed a map from \mathcal{A}_n into \mathcal{A}_n . To proceed, it is known that any bipartite graph has a 2-(node)-coloring. We avoid the wording *coloring* to avoid confusion. Instead, we define *2-Partitions*. Let \mathcal{G}_n^b be the set of all connected undirected bipartite graphs defined over the vertex set \mathcal{I}_n .

Definition 4.6. *Given a graph $G(\mathcal{I}_n, E^b)$ in \mathcal{G}_n^b , a 2-Partition of \mathcal{I}_n with respect to G , is a pair (P_o, P_e) of disjoint subsets of \mathcal{I}_n such that $P_o \cup P_e = \mathcal{I}_n$ and there does not exist an (i, j) in $P_o^2 \cup P_e^2$ such that $ij \in E^b$.*

We eventually restrict G_k to act on the nodes in P_o and P_e separately. For convention, o would refer to *odd* and e to *even*. The dynamics will be presented in such a way, that nodes in P_o (resp. P_e) will be allowed to change actions only at odd (resp. even) time steps. Let us first clearly define a partition of a set.

Definition 4.7. Let X be a set, a partition P_1, \dots, P_m of X is a finite collection of disjoint non-empty subsets of X whose union is X .

The definition to follow serves mainly as a notational clarification, its technical value is rather intuitive.

Definition 4.8. Consider a function f mapping \mathcal{I}_n into some set. Let P_1, \dots, P_m be a partition of \mathcal{I}_n , and let $f|P_l$ be the restriction of f to have domain P_l . Let π be any permutation on $\{1, \dots, m\}$, we consider f to be equal to $(f|P_{\pi(1)}, \dots, f|P_{\pi(m)})$.

Given a 2-Partition, we may ‘decouple’ the dynamics and the following identities would emerge:

Proposition 4.9. Given a pair (G, k) in $\mathcal{G}_n^b \times \mathcal{K}_n$, if we consider a 2-Partition (P_o, P_e) of \mathcal{I}_n with respect to G , then:

1. $G_k a = ((G_k|_{P_o} a) \upharpoonright P_o, (G_k|_{P_e} a) \upharpoonright P_e)$
2. $G_k|_{P_e} G_k|_{P_o} a = (G_k a \upharpoonright P_o, G_k^2 a \upharpoonright P_e)$
3. $G_k^2 a = ((G_k|_{P_e} G_k|_{P_o} a) \upharpoonright P_e, (G_k|_{P_o} G_k|_{P_e} a) \upharpoonright P_o)$.

Proof. The fact that $G_k a = ((G_k|_{P_o} a) \upharpoonright P_o, (G_k|_{P_e} a) \upharpoonright P_e)$ follows from the definition of G_k . We have $(G_k|_{P_o} a) \upharpoonright P_o = (G_k a) \upharpoonright P_o$, so $(G_k|_{P_e} G_k|_{P_o} a) \upharpoonright P_o = (G_k a) \upharpoonright P_o$, and on the other hand since $G_k|_{P_e}$ modifies the action of the players in P_e based only on the actions in P_o , we get:

$$\begin{aligned} (G_k|_{P_e} G_k|_{P_o} a) \upharpoonright P_e &= G_k|_{P_e}((G_k a) \upharpoonright P_o, a \upharpoonright P_e) \upharpoonright P_e \\ &= G_k|_{P_e}((G_k a) \upharpoonright P_o, (G_k a) \upharpoonright P_e) \upharpoonright P_e \\ &= G_k|_{P_e} G_k a \upharpoonright P_e \\ &= G_k^2 a \upharpoonright P_e. \end{aligned}$$

As for the last statement:

$$\begin{aligned} G_k^2 a \upharpoonright P_o &= G_k(G_k a \upharpoonright P_e, G_k a \upharpoonright P_o) \upharpoonright P_o \\ &= G_k|_{P_o}(G_k a \upharpoonright P_e, G_k a \upharpoonright P_o) \upharpoonright P_o \\ &= G_k|_{P_o}(G_k|_{P_e} a \upharpoonright P_e, G_k|_{P_e} a \upharpoonright P_o) \upharpoonright P_o \\ &= (G_k|_{P_o} G_k|_{P_e} a) \upharpoonright P_o. \end{aligned}$$

Similarly, we get $G_k^2 a \upharpoonright P_e = (G_k|_{P_e} G_k|_{P_o} a) \upharpoonright P_o$. □

This fact allows us to say something about the sequence $a, G_k a, G_k^2 a, \dots$ for all a by observing $b, G_k|_{P_e} G_k|_{P_o} b, (G_k|_{P_e} G_k|_{P_o})^2 b, \dots$ for all b . It allows us to observe the process *diagonally*, in a zig-zag fashion. We shall use this fact, and we state the following lemma to formalize the idea.

Proposition 4.10. Given a pair (G, k) in $\mathcal{G}_n^b \times \mathcal{K}_n$, for every cycle C in $\text{CYCLE}_n(G, k)$, if for some action configuration a in C , for any 2-Partition (P_o, P_e) of \mathcal{I}_n with respect to G we have $(G_k|_{P_e} G_k|_{P_o} a) \upharpoonright P_e = a \upharpoonright P_e$ then C has a cardinality of at most 2.

Of course, the condition could have also been written as $(G_k|_{P_o} G_k|_{P_e} a) \upharpoonright P_o = a \upharpoonright P_o$, but that would make no difference in the previous statement: we need the property to hold for all 2-Partitions.

Proof. From Proposition 4.9, we have $G_k^2 a = ((G_k|_{P_e} G_k|_{P_o} a) \upharpoonright P_e, (G_k|_{P_o} G_k|_{P_e} a) \upharpoonright P_o)$. Let C be a cycle in $CYCLE_n(G, k)$, if we have $(G_k|_{P_e} G_k|_{P_o} a) \upharpoonright P_e = a \upharpoonright P_e$ and $(G_k|_{P_o} G_k|_{P_e} a) \upharpoonright P_o = a \upharpoonright P_o$ for some action configuration in C , then $G_k^2 a = a$ and so C has a cardinality of at most 2. \square

To prove that all cycle have cardinality at most 2, we will study for all 2-Partitions (P_o, P_e) and every a in \mathcal{A}_n , the sequence $a, G_k|_{P_e} G_k|_{P_o} a, (G_k|_{P_e} G_k|_{P_o})^2 a, \dots$ and show that this sequence is eventually constant, i.e. there exists a finite time step after which all terms in the sequence become equal. If that is the case, then for any a in \mathcal{A}_n , the sequence $a, G_k^2 a, G_k^4 a, \dots$ is eventually constant and so cycles cannot have a cardinality greater than 2. To this end, we note the following fact:

Proposition 4.11. *Given a pair (G, k) in $\mathcal{G}_n^b \times \mathcal{K}_n$ and a 2-Partition (P_o, P_e) of \mathcal{I}_n , for every player i in \mathcal{I}_n , if $c \in \{\mathbb{W}, \mathbb{B}\}$ is a strong assignment for player i , then for every a in \mathcal{A}_n , if $(G_k|_{P_e} G_k|_{P_o} a)_i = c$, then $((G_k|_{P_e} G_k|_{P_o})^m a)_i = c$ for all positive integers m .*

Proof. If i belongs to P_e , we have $(G_k|_{P_e} G_k|_{P_o} a)_i = (G_k^2 a)_i$, and the result then follows from the definition of strong assignment. If i belongs to P_o , then $(G_k|_{P_e} G_k|_{P_o} a)_i = (G_k|_{P_o} a)_i$. So if $(G_k|_{P_e} G_k|_{P_o} a)_i = c$ then $(G_k|_{P_o} a)_i = c$ and since:

$$\begin{aligned} ((G_k|_{P_e} G_k|_{P_o})^m a)_i &= (G_k|_{P_e} (G_k|_{P_o} G_k|_{P_e})^{m-1} G_k|_{P_o} a)_i \\ &= ((G_k|_{P_o} G_k|_{P_e})^{m-1} G_k|_{P_o} a)_i \\ &= (G_k^{2m-2} G_k|_{P_o} a)_i, \end{aligned}$$

the result would follow from the definition of strong assignment. \square

We present a rather natural statement that will be used in the inductive argument to prove the theorem. Although the statement is intuitively simple, we give an effort to state it carefully so to take care of any minor technicalities involved. We first give a simplistic informal version of it for general graphs: given a triplet (G, k, a) in $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$, let us consider the sequence $a, G_k a, G_k a^2, \dots$. If some player *never* changes action along this sequence, then we may delete the player from the graph and modify the thresholds of the neighboring players in such a way that the effect of the action played by the player is ‘seen’ by the neighbors. We do so by keeping the thresholds of the neighbors unchanged if that action is \mathbb{W} , and by decreasing the thresholds by 1 (to a minimum of zero) if that action is \mathbb{B} . In this case, we would obtain a different triplet (G', k', a') where all the players in $V(G')$ play the exact actions as in the sequence $a, G_k a, G_k a^2, \dots$. We now specialize the previous idea to meet our needs for the bipartite case.

Proposition 4.12. *Let (G, k) be a pair in $\mathcal{G}_n^b \times \mathcal{K}_n$, (P_o, P_e) be a 2-Partition of \mathcal{I}_n and a an action configuration in \mathcal{A}_n . If there exists a player i and an action c , such that $((G_k|_{P_e} G_k|_{P_o})^m a)_i = c$ for all non-negative integers m , consider H to be the induced subgraph of G over $\mathcal{I}_n \setminus \{i\}$. Suppose G' is a connected component of H with vertex set \mathcal{J} , define a' to be the action configuration a restricted to the players in \mathcal{J} , define P'_o and P'_e to be*

$P_o \cap \mathcal{J}$ and $P_e \cap \mathcal{J}$ respectively, and define k' to be the map from \mathcal{J} into \mathbb{N} such that $k' = k$ on $\mathcal{J} \setminus \mathcal{N}_i$, $k' = k$ on $\mathcal{N}_i \cap \mathcal{J}$ if $c = \mathbb{W}$ and $k' = (k - 1) \vee 0$ on $\mathcal{N}_i \cap \mathcal{J}$ if $c = \mathbb{B}$. Then:

$$((G_k|_{P_e} G_k|_{P_o})^m a)_j = ((G'_{k'}|_{P'_e} G'_{k'}|_{P'_o})^m a')_j,$$

for all non-negative integers m and all players j in \mathcal{J} .

Proof. It would be enough to show that the local decision rules of the players in $\mathcal{J} \cap \mathcal{N}_i$ does not change, simply that

$$(G_k|_{P_e} G_k|_{P_o} a)_j = (G'_{k'}|_{P'_e} G'_{k'}|_{P'_o} a')_j.$$

Let j be a player in $\mathcal{J} \cap \mathcal{N}_i$, either j belongs to P_o or j belongs to P_e . Suppose j belongs to P_o , then $(G_k|_{P_e} G_k|_{P_o} a)_j = (G_k|_{P_o} a)_j$, and $(G_k|_{P_o} a)_j = \mathbb{B}$ if and only if at least k_j nodes in \mathcal{N}_j play \mathbb{B} in a , or equivalently at least k'_j nodes in $\mathcal{N}_j \setminus \{i\}$ play \mathbb{B} in a' , since k' takes into account the action of player i . And that is equivalent to $(G'_{k'}|_{P'_e} G'_{k'}|_{P'_o} a')_j = \mathbb{B}$. Similarly, if j belongs to P_e we get that $(G_k|_{P_e} G_k|_{P_o} a)_j = (G'_{k'}|_{P'_e} (G_k|_{P_o} a \upharpoonright \mathcal{J}))_j = (G'_{k'}|_{P'_e} G'_{k'}|_{P'_o} a)_j$. \square

With this settled, we go now to proving our theorem.

Proposition 4.13. *For any positive integer n , any tree T in \mathcal{G}_n and any threshold distribution k in \mathcal{K}_n , every cycle C in $\text{CYCLE}_n(T, k)$ has cardinality less than or equal to 2.*

Proof. We prove this statement by induction. The statement can be exhaustively checked for $n = 2$, and it is trivial for $n = 1$. So let us assume that it holds when the number of players is less than or equal to $n \geq 3$. Let (T, k, a) be a triplet in $\mathcal{G}_{n+1}^b \times \mathcal{K}_{n+1} \times \mathcal{A}_{n+1}$ where T is a tree and consider a 2-Partition (P_o, P_e) of \mathcal{I}_{n+1} . We consider the sequence $a, G_k|_{P_e} G_k|_{P_o} a, (G_k|_{P_e} G_k|_{P_o})^2 a, \dots$. We claim that at least one node will never change colors along the sequence after some finite time step. If there exists a non-valid node, then clearly that is the case. If all nodes are valid, then (since $n \geq 2$) necessarily there exists a parent and a leaf. By Lemma 4.4, the parent has a strong assignment, and so it stops changing colors along the sequence after some finite time step. Either way, we may then apply Proposition 4.12 on that node at the time step when the color becomes constant, removing the node and updating the thresholds accordingly. The result follows since it holds on the obtained connected components by assumption. \square

We proceed to extend such results to general graphs. Unfortunately, the notion of strong assignment does not extend to the general case, we construct a different approach.

4.3 General Graphs

Given a pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ we will have this pair undergo two procedures: a bipartite-expansion and a symmetric-expansion. The bipartite-expansion will produce a pair (G', k') where G' is bipartite, and the symmetric-expansion will produce a pair (G'', k'') where G'' has only nodes with odd degrees and k'' induces a majority decision rule. We show that by restricting the set of initial action profiles on (G', k') and (G'', k'') , these two pairs simulate (G, k) (see Lemmas 4.20, 4.21 and 4.22). Given a pair (G, k) , we are then able to assume without any loss of generality (in proving the bound) that G is bipartite, has only odd-degree

nodes and that k induces a majority decision rule. The bipartite property allows us to consider a 2-Partition of the graph and decouple the parallel decision scheme into a sequential process (with respect to the 2-Partition) as was done in the case of Trees. The majority decision rule then allows us to prove (by counting specific edges in the graph) that no node can flip infinitely many times along this sequential process. This argument establishes fixed-point convergence of the sequential process, translating to an upper-bound of 2 on the length of the convergence cycles for the parallel dynamics.

For technical and notational convenience, we will assume (throughout the rest of the section) that when dealing with only two graphs G and G' in \mathcal{G}_n and $\mathcal{G}_{n'}$ respectively with $n < n'$, the set $V(G)$ lies inside $V(G')$, i.e. \mathcal{I}_n lies inside $\mathcal{I}'_{n'}$. This only removes the need to explicitly identify a subset of $V(G')$ to $V(G)$.

First, recall that we denote by \mathcal{G}_n^b the set of all connected undirected bipartite graphs defined over the vertex set \mathcal{I}_n . We begin by formally introducing the bipartite-expansion and the symmetric-expansion.

Definition 4.14. *Given a pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ where G is non-bipartite, we construct a pair (G', k') in $\mathcal{G}_{2n}^b \times \mathcal{K}_{2n}$ as outlined in the procedure to follow. We refer to (G', k') as the bipartite-expansion of (G, k) .*

Bipartite-expansion procedure — Given a pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ where G is non-bipartite and has edge set E , we construct a pair (G', k') in $\mathcal{G}_{2n}^b \times \mathcal{K}_{2n}$ as follows. We suppose G' is equal to (\mathcal{I}_{2n}, E') , and partition \mathcal{I}_{2n} into two sets \mathcal{I}_n and \mathcal{J} . We define a bijection ϕ from \mathcal{J} into \mathcal{I}_n and define E' to be the set of undirected edges on \mathcal{I}_{2n} such that for i and j in \mathcal{I}_{2n} , $\{i, j\} \in E'$ if and only if $(i, j) \in \mathcal{I}_n \times \mathcal{J}$ and $\{i, \phi(j)\} \in E$. Finally, set k' to be equal to k on \mathcal{I}_n and $k \circ \phi$ on \mathcal{J} .

Example 4.15. *Under a bipartite-expansion, a cycle over 3 nodes (left) becomes a cycle over 6 nodes (right).*




The integer tag on each node represents its corresponding threshold.

We define \mathcal{S}_n to be the set of all pairs (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ such for each player i in \mathcal{I}_n , the degree d_i is odd and k_i is equal to $(d_i + 1)/2$. We refer to \mathcal{S}_n as the set of *symmetric models*, in the sense that for (G, k) in \mathcal{S}_n the property is such that for any action profile a in \mathcal{A}_n , and any player i , the action $(G_k a)_i$ is the action played by the majority in \mathcal{N}_i with respect to the action profile a . In this case, the two actions \mathbb{B} and \mathbb{W} are treated as having equal weights by all players in the network.

Definition 4.16. Given a pair (G, k) in $(\mathcal{G}_n \times \mathcal{K}_n) \setminus \mathcal{S}_n$, we construct a pair (G', k') in $\mathcal{G}'_n \times \mathcal{K}'_n$ as outline in the procedure to follow. We refer to (G', k') as a one-step symmetric-expansion of (G, k) .

One-step symmetric-expansion procedure — Given a pair (G, k) in $(\mathcal{G}_n \times \mathcal{K}_n) \setminus \mathcal{S}_n$, we construct a pair (G', k') in $\mathcal{G}'_n \times \mathcal{K}'_n$ as follows. We suppose that G is equal to (\mathcal{I}_n, E) , and choose a player i in \mathcal{I}_n such that either d_i is even, or d_i is odd and k_i is not equal to $(d_i + 1)/2$. Surely such a node exists since (G, k) does not belong to \mathcal{S}_n . We call the node i the **pivot** node in the one-step symmetric-expansion of (G, k) into (G', k') . We construct an instance (G', k') in $\mathcal{G}_{n+3(d_i+1)} \times \mathcal{K}_{n+3(d_i+1)}$. We suppose that G' is equal to $(\mathcal{I}_{n+3(d_i+1)}, E')$ and partition $\mathcal{I}_{n+3(d_i+1)}$ into $\mathcal{I}_n, P_1, \dots, P_{d_i+1}$ where each partition different than \mathcal{I}_n has cardinality exactly equal 3. We define E' to be the undirected set of edges such that E' contains E , and for every m , we suppose $P_m = \{j, j', j''\}$ and let E' contain jj', jj'' and ij . To visualize the obtained graph structure G' , we attached $d_i + 1$ three-node Y-shaped graphs to node i . Finally, we set k' to be equal to k on $\mathcal{I}_n \setminus \{i\}$, to be equal to $d_i + 1$ at i , equal to 2 on the remaining nodes having degree 3 and equal to 1 everywhere else.

Example 4.17. The figure below shows a one-step symmetric expansion (right) of a pair (G, k) (left) where G is a cycle graph over 3 nodes.



The integer tag on each node represents its corresponding threshold. The pivot in this example is a node (in the cycle graph) having a threshold of 1.

Iterated one-step symmetric-expansion of pair in $(\mathcal{G}_n \times \mathcal{K}_n) \setminus \mathcal{S}_n$ eventually yields a pair in \mathcal{S}_n . We formalize the idea:

Definition 4.18. Given a pair (G_0, k_0) in $(\mathcal{G}_n \times \mathcal{K}_n) \setminus \mathcal{S}_n$, we construct a finite sequence $(G_0, k_0), (G_1, k_1), \dots, (G_m, k_m)$ for some positive integer m , where (G_l, k_l) is a one-step symmetric-expansion of (G_{l-1}, k_{l-1}) and (G_m, k_m) belongs to $\mathcal{S}_{n'}$ for some n' . We refer to (G_m, k_m) as the symmetric-expansion of (G, k) .

We note that the symmetric-expansion of (G, k) is uniquely defined (up to isomorphism¹³) regardless of the order the pivot nodes were chosen. However, we shall not write a formal

¹³ Two pairs (G, k) and (G', k') in $\mathcal{G}_n \times \mathcal{K}_n$ are said to be isomorphic if there exists a bijective map ϕ from $V(G)$ into $V(G')$ such that $ij \in E(G)$ if and only if $\phi_i \phi_j \in E(G')$ and $k_i = (k' \circ \phi)_i$ for i in $V(G)$.

proof for this fact, it would suffice to say that when we perform the one-step symmetric expansion of a pair (G, k) , we leave all neighborhoods and thresholds unchanged for all nodes different than the pivot.¹⁴

Our goal is to prove a bound (equal to 2) on the length of convergence cycles for all pairs (G, k) . We first establish that such a bound applies for a pair (G, k) if it applies for either its symmetric-expansion (Lemmas 4.20 and 4.21) or its bipartite-expansion (Lemma 4.22). We combine both expansions (Proposition 4.23), and thus show that, to get the general result, we need only prove the bound for symmetric models (G, k) where G is bipartite (Lemmas 4.24 and 4.25).

Definition 4.19. *We define the set \mathcal{M} to be a subset of $\bigcup_{n \geq 1} \mathcal{G}_n \times \mathcal{K}_n$ such that (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ belongs to \mathcal{M} if and only if for every C in $\text{CYCLE}_n(G, k)$, the cardinality of C is less than or equal to 2.*

What we ultimately show is that $\mathcal{M} = \bigcup_{n \geq 1} \mathcal{G}_n \times \mathcal{K}_n$.

Lemma 4.20. *Given a pair (G, k) in $(\mathcal{G}_n \times \mathcal{K}_n) \setminus \mathcal{S}_n$, define (G', k') to be the one-step symmetric-expansion of (G, k) . If (G', k') belongs to \mathcal{M} then (G, k) belongs to \mathcal{M} .*

Proof. Let node i be the *pivot* node in the one-step symmetric-expansion of (G, k) into (G', k') . We suppose G' belongs to $\mathcal{G}_{n'}$, and consider the induced subgraph H in G' over the vertex set $\mathcal{I}_{n'} \setminus \mathcal{I}_n$. The subgraph H necessarily consists (by construction of G') of $d_i + 1$ connected components, each consisting of three vertices. We set integers b_i and w_i to be equal to k_i and $d_i - k_i + 1$ respectively, we then partition $\mathcal{I}_{n'} \setminus \mathcal{I}_n$ into $W_1, \dots, W_{b_i}, B_1, \dots, B_{w_i}$ such that each partition contains the set of nodes in one of the connected components. Suppose that (G', k') belongs to \mathcal{M} , we show that (G, k) belongs to \mathcal{M} . Let C be an element of $\text{CYCLE}_n(G, k)$ and suppose C has cardinality greater or equal to 3. In particular, suppose C is equal to $\{a_1, \dots, a_m\}$ for $m \geq 3$ where $a_{l+1} = G_k a_l$ for $1 \leq l < m$ and $a_1 = G_k a_m$.

We define a map α from \mathcal{A}_n into $\mathcal{A}_{n'}$ such that for a in \mathcal{A}_n ,

$$\alpha a = \begin{cases} a & \text{on } \mathcal{I}_n \\ \mathbb{B} & \text{on } B_1 \cup \dots \cup B_{w_i} \\ \mathbb{W} & \text{on } W_1 \cup \dots \cup W_{b_i} \end{cases} .$$

First, the map α is clearly injective. Then, $\alpha a_1, \dots, \alpha a_m$ are distinct elements of $\mathcal{A}_{n'}$. Second, it can be checked that

$$\alpha(G_k a) = G'_{k'}(\alpha a).$$

To see this fact, notice that for every b_1 and b_2 in \mathcal{A}_n every node in $\mathcal{I}_{n'} \setminus \mathcal{I}_n$ has the same color both in αb_1 and αb_2 , and the same color both in αb_1 and $G'_{k'} \alpha b_1$. Since every node in \mathcal{I}_n other than the pivot i keeps the same neighborhood and threshold, we need only show that $(\alpha G_k a)_i = (G'_{k'} \alpha a)_i$. To this end, node i is \mathbb{B} in $G_k a$ if and only if at least k_i neighbors

¹⁴ Given a pair (G, k) in $(\mathcal{G}_n \times \mathcal{K}_n) \setminus \mathcal{S}_n$, it can be easily checked that if (G', k') be the one-step symmetric-expansion of (G, k) , then G is bipartite if and only if G' is bipartite. Similarly, given a pair (G, k) in $(\mathcal{G}_n \setminus \mathcal{G}_n^b) \times \mathcal{K}_n$, let (G', k') be the bipartite-expansion of (G, k) , then (G, k) belongs to \mathcal{S}_n if and only if (G', k') belongs to \mathcal{S}_{2n} .

of i in G are \mathbb{B} in a , and that is the case if and only if at least $k_i + w_i$ neighbors of i in G' are \mathbb{B} in αa , or equivalently at least $d^G(i) + 1 = (d^{G'}(i) + 1)/2$ neighbors of i in G' are \mathbb{B} in αa . Finally, it follows that $\alpha C = \{\alpha a_1, \dots, \alpha a_m\}$ is a cycle in $CYCLE_{n'}(G', k')$ contradicting the fact that $m \geq 3$. \square

Lemma 4.21. *Given a pair (G, k) in $(\mathcal{G}_n \times \mathcal{K}_n) \setminus \mathcal{S}_n$, define (G', k') to be the symmetric-expansion of (G, k) . If (G', k') belongs to \mathcal{M} then (G, k) belongs to \mathcal{M} .*

Proof. Given that the pair (G, k) belongs to $(\mathcal{G}_n \times \mathcal{K}_n) \setminus \mathcal{S}_n$, we can construct a finite sequence $(G_1, k_1), \dots, (G_{m-1}, k_{m-1})$ in such a way that (G_1, k_1) is the one-step symmetric expansion of (G, k) , (G', k') is the one-step symmetric-expansion of (G_{m-1}, k_{m-1}) and (G_l, k_l) is the one-step symmetric-expansion of (G_{l-1}, k_{l-1}) for $1 < l < m$. If (G', k') belongs to \mathcal{M} , then (G_{m-1}, k_{m-1}) belongs to \mathcal{M} by Lemma 4.20. Recursively, it follows that (G, k) belongs to \mathcal{M} . \square

Lemma 4.22. *Given a pair (G, k) in $(\mathcal{G}_n \setminus \mathcal{G}_n^b) \times \mathcal{K}_n$, define (G', k') to be the bipartite-expansion of (G, k) . If (G', k') belongs to \mathcal{M} then (G, k) belongs to \mathcal{M} .*

Proof. The graph G' is bipartite and has vertex set \mathcal{I}_{2n} , let us partition \mathcal{I}_{2n} into \mathcal{I}_n and \mathcal{J} , then $(\mathcal{I}_n, \mathcal{J})$ forms a 2-Partition with respect to G' by construction. We define a bijection ϕ from \mathcal{J} into \mathcal{I}_n such that for j_1 and j_2 in \mathcal{J} , $j_1\phi(j_2) \in E'$ if and only if $\phi(j_1)\phi(j_2) \in E$. Given an action configuration in \mathcal{A}_n , we define the map α from \mathcal{A}_n into \mathcal{A}_{2n} such that for a in \mathcal{A}_n we have:

$$\alpha a = (a, a \circ \phi).$$

It then follows that:

$$G'_{k'}(\alpha a) = \alpha(G_k a).$$

The map α is clearly injective, it then follows that if $C = \{a_1, \dots, a_m\}$ is a cycle in $CYCLE_n(G, k)$ with $m > 2$, then $\alpha C = \{\alpha a_1, \dots, \alpha a_m\}$ is a cycle in $CYCLE_n(G', k')$ contradicting the fact that (G', k') belongs to \mathcal{M} . \square

Proposition 4.23. *Given a pair (G, k) in $((\mathcal{G}_n \setminus \mathcal{G}_n^b) \times \mathcal{K}_n) \setminus \mathcal{S}_n$, we define (G_1, k_1) and (G_2, k_2) to be respectively the bipartite-expansion and the symmetric expansion of (G, k) . If (G'_1, k'_1) is the symmetric-expansion of (G_1, k_1) and (G'_2, k'_2) is the bipartite-expansion of (G_2, k_2) , then (G'_1, k'_1) is equal to (G'_2, k'_2) (up to isomorphism¹⁵).*

Proof. It suffices to prove existence of a bijective map ϕ from $V(G'_2)$ to $V(G'_1)$, such that $ij \in E(G'_1)$ if and only if $\phi_i\phi_j \in E(G'_2)$ and $(k'_1)_i = (k'_2 \circ \phi)_i$ for i in $V(G'_1)$. Such a map could be easily constructed following the expansion procedure. We omit the construction. \square

Lemma 4.24. *Given a pair (G, k) in $((\mathcal{G}_n \setminus \mathcal{G}_n^b) \times \mathcal{K}_n) \setminus \mathcal{S}_n$, define (G', k') to be the bipartite-expansion of (G, k) and (G'', k'') to be the symmetric-expansion of (G', k') . If (G'', k'') belongs to \mathcal{M} then (G, k) belongs to \mathcal{M} .*

Proof. If (G'', k'') belongs to \mathcal{M} , then (G', k') belongs to \mathcal{M} by Lemma 4.21. It follows that (G, k) belongs to \mathcal{M} by Lemma 4.22. \square

¹⁵ See footnote 13.

Lemma 4.25. *The set \mathcal{M} is equal to $\bigcup_{n \geq 1} \mathcal{G}_n \times \mathcal{K}_n$ if and only if \mathcal{M} contains $\mathcal{S}_n \cap (G_n^b \times \mathcal{K}_n)$ for every positive integer n .*

Proof. It is clear that if $\mathcal{M} = \bigcup_{n \geq 1} \mathcal{G}_n \times \mathcal{K}_n$ then $\mathcal{S}_n \cap (G_n^b \times \mathcal{K}_n) \subset \mathcal{M}$. To prove the converse, given (G, k) an element of $\bigcup_{n \geq 1} \mathcal{G}_n \times \mathcal{K}_n$, we define (G', k') to be the bipartite-expansion of (G, k) and (G'', k'') to be the symmetric-expansion of (G', k') . We have that (G'', k'') belongs to $\mathcal{S}_n \cap (G_n^b \times \mathcal{K}_n)$. If \mathcal{M} contains $\mathcal{S}_n \cap (G_n^b \times \mathcal{K}_n)$, it then follows by Lemma 4.24 that \mathcal{M} contains (G, k) . \square

We proceed to show that \mathcal{M} contains $\mathcal{S}_n \cap (G_n^b \times \mathcal{K}_n)$ for every positive integer n .

Lemma 4.26. *Given a pair (G, k) in $G_n^b \times \mathcal{K}_n$, we consider a 2-Partition (P_o, P_e) of \mathcal{I}_n with respect to G . Then (G, k) belongs to \mathcal{M} if and only if for every a in \mathcal{A}_n , there exists some integer T , such that $(G_k|_{P_e} G_k|_{P_o})^T a = (G_k|_{P_e} G_k|_{P_o})^{T+1} a$.*

Proof. The statement of this lemma is equivalent to the statement of Proposition 4.10. \square

Definition 4.27 (Conflict Link). *Given (G, a) in $\mathcal{G}_n \times \mathcal{A}_n$ with $G = (\mathcal{I}_n, E)$, we call a conflict link in G with respect to a , an element ij of E such that a_i and a_j are not equal. We denote by $E_c^G(a)$ the set of all conflict links in G with respect to a .*

Lemma 4.28. *The set \mathcal{M} contains $\mathcal{S}_n \cap (G_n^b \times \mathcal{K}_n)$ for every positive integer n .*

Proof. Let (G, k) in $\mathcal{S}_n \cap (G_n^b \times \mathcal{K}_n)$ be given and consider a 2-Partition (P_o, P_e) of \mathcal{I}_n . Let a be an action profile in \mathcal{A}_n . By Lemma 4.26 it would be enough to show that $(G_k|_{P_e} G_k|_{P_o})^T a = (G_k|_{P_e} G_k|_{P_o})^{T+1} a$ for some non-negative integer T . In that case, it would be enough to prove that for every b in \mathcal{A}_n ,

$$G_k|_{P_o} b \neq b \quad \text{if and only if} \quad |E_c^G(G_k|_{P_o} b)| < |E_c^G(b)|,$$

and that similarly,

$$G_k|_{P_e} b \neq b \quad \text{if and only if} \quad |E_c^G(G_k|_{P_e} b)| < |E_c^G(b)|.$$

To show that, we state the following: for node i in \mathcal{I}_n , $a_i \neq (G_k a)_i$ if and only if the majority of the players in \mathcal{N}_i are not playing a_i , or equivalently, if and only if i can decrease the number of conflict links by switching action. \square

Theorem 4.29. *For every positive integer n , every (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ and every C in $CYCLE_n(G, k)$, the cardinality of C is less than or equal to 2.*

Proof. The statement of the theorem is equivalent to $\mathcal{M} = \bigcup_{n \geq 1} \mathcal{G}_n \times \mathcal{K}_n$. The fact that $\mathcal{M} = \bigcup_{n \geq 1} \mathcal{G}_n \times \mathcal{K}_n$ follows directly from Lemma 4.28 and Lemma 4.25. \square

We will proceed to extend the result to the extension model. But before that, we perform a preparatory step to provide more insight on the problem. The primary model exhibits a (positive) monotonicity property in the following sense: a node is more likely to play \mathbb{B} as the number of \mathbb{B} playing neighbors increases. In what follows, we invert this property.

4.4 Inverted Rules

For a positive integer n , let (G, k) be an element of $\mathcal{G}_n \times \mathcal{K}_n$. We define the map $\neg G_k$ from \mathcal{A}_n into \mathcal{A}_n such that for $a \in \mathcal{A}_n$, $(\neg G_k a)_i \neq (G_k a)_i$ for all i in \mathcal{I}_n . Then $(\neg G_k a)_i$ is equal to \mathbb{B} if and only if *at most* $k_i - 1$ players are in $a^{-1}(\mathbb{B}) \cap \mathcal{N}_i$. Likewise, $(\neg G_k a)_i$ is equal to \mathbb{B} if and only if *at least* $d_i - k_i + 1$ players are in $a^{-1}(\mathbb{W}) \cap \mathcal{N}_i$. If we extend the definition of $CYCLE_n(G, k)$ to $\neg G_k$, and denote the set of cycles by $CYCLE_n(\neg G_k)$, we get the following theorem:

Theorem 4.30. *For every positive integer n , every (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ and every C in $CYCLE_n(\neg G_k)$, the cardinality of C is less than or equal to 2.*

Proof. For a positive integer n , let (G, k) be an element of $\mathcal{G}_n \times \mathcal{K}_n$. We will simulate $\neg G_k$ using an instance of the primary model. For i in \mathcal{I}_n , $(\neg G_k a)_i = \mathbb{B}$ if and only if at most $k_i - 1$ players are in $a^{-1}(\mathbb{B}) \cap \mathcal{N}_i$. Since the bipartite-expansion procedure is only defined for non-bipartite graphs, for notational (and technical) convenience, we will assume the graph G to be non-bipartite. The bipartite case may be easily recreated from this proof by considering a 2-Partition directly. Let (G', k') be the bipartite expansion of (G, k) , then G' is bipartite, has vertex set \mathcal{I}_{2n} and edge set E' . Let us partition \mathcal{I}_{2n} into \mathcal{I}_n and \mathcal{J} . We define a bijection ϕ from \mathcal{J} into \mathcal{I}_n such that for j_1 and j_2 in \mathcal{J} , $j_1 \phi(j_2) \in E'$ if and only if $\phi(j_1) \phi(j_2) \in E$. We define \bar{k} such that \bar{k}_i is equal to k'_i for i in \mathcal{J} and equal to $d_i - k'_i + 1$ for i in \mathcal{I}_n . For a in \mathcal{A}_n , we first define $\neg a$ to be the element of \mathcal{A}_n such that $\neg a_i \neq a_i$ for all i in \mathcal{I}_n , and then define α to be the map from \mathcal{A}_n into \mathcal{A}_{2n} such that for a in \mathcal{A}_n :

$$\alpha a = (a, (\neg a) \circ \phi).$$

This done, we get:

$$G'_{\bar{k}} \alpha a = G'_{\bar{k}}(a, (\neg a) \circ \phi) = (\neg G_k a, (G_k a) \circ \phi).$$

To see this, for i in \mathcal{I}_n , $(G'_{\bar{k}} \alpha a)_i$ is \mathbb{B} if and only if at least $d_i - k_i + 1$ players are in $\phi^{-1}((\neg a)^{-1}(\mathbb{B})) \cap \mathcal{N}_{G'}(i)$. Or equivalently, at least $d_i - k_i + 1$ are in $\phi^{-1}(a^{-1}(\mathbb{W})) \cap \mathcal{N}_{G'}(i)$. But at least $d_i - k_i + 1$ are in $\phi^{-1}(a^{-1}(\mathbb{W})) \cap \mathcal{N}_{G'}(i)$ if and only if at least $d_i - k_i + 1$ are in $a^{-1}(\mathbb{W}) \cap \mathcal{N}_G(i)$ since $\mathcal{N}_G(i) = \phi(\mathcal{N}_{G'}(i))$. It then follows that $(G'_{\bar{k}} \alpha a)_i$ is \mathbb{B} if and only if at most $k_i - 1$ players are in $a^{-1}(\mathbb{B}) \cap \mathcal{N}_G(i)$.

Let $C = \{a_1, \dots, a_m\}$ be a cycle in $CYCLE_n(\neg G_k)$, then $\neg C = \{\neg a_1, \dots, \neg a_m\}$ is a cycle in $CYCLE_n(G_k)$ and so $\{(a_1, (\neg a_1) \circ \phi), \dots, (a_m, (\neg a_m) \circ \phi)\}$ is a cycle in $CYCLE_{2n}(G_{\bar{k}})$. Therefore, m is less than or equal to 2. \square

This ‘inverted’ rule may be viewed as having negative weights equal to -1 along all edges then making the thresholds negative. We proceed to study the extension model.

4.5 Convergence Cycles for the Extension Model

We extend the convergence cycle result to the extension model in three steps. We initially assume that no self loops are present. We extend the result first to the case where weights are allowed to be equal to either 1 or -1 and then to arbitrary weights. We finally allow weighted self-loops to be present, and prove the theorem.

We begin by assuming that no self-loops are present. We then expand \mathcal{G}_n to incorporate the weights. We define \mathcal{W}_n to be the set of connected graph over the vertex set \mathcal{I}_n with weighted edges. The set \mathcal{W}_n is then isomorphic to the set of symmetric matrices in $\mathbb{R}^{n \times n}$ having zero elements along the diagonal. Given a weighted graph W of \mathcal{W}_n , we denote the set of edges by $E(W)$, the weight on edge ij in $E(W)$ is then denoted by w_{ij} . As a natural extension of previous definitions, given a pair (W, q) we define the map W_q from \mathcal{A}_n into \mathcal{A}_n such that $(W_q a)_i = \mathbb{B}$ if and only if $\sum_{j \in \mathcal{N}_i} w_{ij} \mathbf{1}_{\{\mathbb{B}\}}(a_j) \geq q_i \sum_{j \in \mathcal{N}_i} w_{ij}$.

We define $\hat{\mathcal{K}}_n$ to be the set of maps from \mathcal{I}_n into the set of integers \mathbb{Z} . Following a similar reason as in Section 3, without any loss in generality, we substitute the set \mathcal{Q}_n with the set $\hat{\mathcal{K}}_n$ as defined; the thresholds can be negative in the extension model. For a pair (W, k) in $\mathcal{W}_n \times \hat{\mathcal{K}}_n$, the map W_k extends naturally from W_q for the pair (W, q) in $\mathcal{W}_n \times \mathcal{Q}_n$, in the way that for an action configuration a in \mathcal{A}_n , $(W_k a)_i = \mathbb{B}$ if and only if $\sum_{j \in \mathcal{N}_i} w_{ij} \mathbf{1}_{\{\mathbb{B}\}}(a_j) \geq k_i$. Similarly, we define $CYCLE_n(W, k)$ to extends naturally from Section 3.

We now restrict the weights to be in $\{-1, 1\}$. Let (W, k) be a pair in $\mathcal{W}_n \times \hat{\mathcal{K}}_n$, we define, for each i in \mathcal{I}_n , \mathcal{N}_i^+ and \mathcal{N}_i^- to be a partition of \mathcal{N}_i such that w_{ij} is positive only if j belongs \mathcal{N}_i^+ . It then follows that $(W_k a)_i$ is \mathbb{B} if and only if $|a^{-1}(\mathbb{B}) \cap \mathcal{N}_i^+| - |a^{-1}(\mathbb{B}) \cap \mathcal{N}_i^-| \geq k_i$.

Theorem 4.31. *For any positive integer n , and any pair (W, k) in $\mathcal{W}_n \times \hat{\mathcal{K}}_n$ where all weights are in $\{-1, 1\}$, every cycle C in $CYCLE_n(W, k)$ has cardinality less than or equal to 2.*

Proof. Given a pair (W, k) in $\mathcal{W}_n \times \hat{\mathcal{K}}_n$ where all weights are in $\{-1, 1\}$, we construct a graph G' over \mathcal{I}_{2n} with edge set E' as follows. We partition \mathcal{I}_{2n} into \mathcal{I}_n and \mathcal{J} , and set up a bijection ϕ from \mathcal{J} into \mathcal{I}_n . We assume that at least two weights are of opposite sign, otherwise nothing needs to be done. We let E^+ and E^- be a partition of $E(W)$ such that the edges in E^+ and E^- have positive and negative weights respectively. We then let E' contain both ij and $\phi(i)\phi(j)$ for all ij in E^+ , and $i\phi(j)$ for all ij in E^- . Given a node i , we denote by d_i^+ and d_i^- to be the number of edges having i as an endpoint in E^+ and E^- respectively. In this sense, $d_i = d_i^+ + d_i^-$. We construct k' to be an element of $\hat{\mathcal{K}}_{2n}$ such that k'_i equals $k_i + d_i^-$ for i in \mathcal{I}_n and equals $d_i^+ - k_i + 1$ for i in \mathcal{J} . Without any loss of generality, we assume that all nodes are valid. Therefore, $-d_i^- \leq k_i \leq d_i^+$ for all i in \mathcal{I}_n .

For a in \mathcal{A}_n , we first define $\neg a$ to be the element of \mathcal{A}_n such that $\neg a_i \neq a_i$ for all i in \mathcal{I}_n , and then define α to be the map from \mathcal{A}_n into \mathcal{A}_{2n} such that for a in \mathcal{A}_n :

$$\alpha a = (a, (\neg a) \circ \phi).$$

This done:

$$G'_{k'} \alpha a = G'_{k'} (a, (\neg a) \circ \phi) = (W_k a, (\neg W_k a) \circ \phi).$$

To see this, for i in \mathcal{I}_n , $(G'_{k'} \alpha a)_i = \mathbb{B}$ if and only if at least $k_i + d_i^-$ players are in $(\alpha a)^{-1}(\mathbb{B}) \cap \mathcal{N}_i^{G'}$, or equivalently:

$$|a^{-1}(\mathbb{B}) \cap \mathcal{N}_i^+| + |a^{-1}(\mathbb{W}) \cap \mathcal{N}_i^-| \geq k_i + d_i^-.$$

But $|a^{-1}(\mathbb{W}) \cap \mathcal{N}_i^-| = d_i^- - |a^{-1}(\mathbb{B}) \cap \mathcal{N}_i^-|$. Therefore, $(G'_{k'} \alpha a)_i = \mathbb{B}$ if and only if $(W_k a)_i = \mathbb{B}$.

Similarly, for i in \mathcal{J} , $(G'_{k'} \alpha a)_i = \mathbb{B}$ if and only if at least $d_i^+ - k_i + 1$ players are in $(\alpha a)^{-1}(\mathbb{B}) \cap \mathcal{N}_i^{G'}$, or equivalently

$$|a^{-1}(\mathbb{W}) \cap \mathcal{N}_i^+| + |a^{-1}(\mathbb{B}) \cap \mathcal{N}_i^-| \geq d_i^+ - k_i + 1.$$

The previous proposition is also equivalent to:

$$d_i^+ - |\neg a^{-1}(\mathbb{B}) \cap \mathcal{N}_i^+| - |a^{-1}(\mathbb{B}) \cap \mathcal{N}_i^-| < k_i.$$

But, $d_i^+ - |\neg a^{-1}(\mathbb{B}) \cap \mathcal{N}_i^+| = |a^{-1}(\mathbb{B}) \cap \mathcal{N}_i^+|$. It follows that $(G'_{k'} \alpha a)_i = (\neg W_k a)_i$.

Let $C = \{a_1, \dots, a_m\}$ be a cycle in $CYCLE_n(W, k)$, then $\{(a_1, (\neg a_1) \circ \phi), \dots, (a_m, (\neg a_m) \circ \phi)\}$ is a cycle in $CYCLE_{2n}(G'_{k'})$. Therefore, m is less than or equal to 2. \square

We now extend the result to the extension model allowing non-equal weights on edges but no self-loops. We put no restrictions on the weights in \mathcal{W}_n other than being non-zero reals. Without any loss of generality, we may assume the weights to be rational (see [18]). In this case, we can multiply all weights by a common factor and have them be integers. This said, without any loss of generality, we assume the weights to be integers.

Theorem 4.32. *For any positive integer n , and any pair (W, k) in $\mathcal{W}_n \times \hat{\mathcal{K}}_n$, every cycle C in $CYCLE_n(W, k)$ has cardinality less than or equal to 2.*

Proof. Given a pair (W, k) in $\mathcal{W}_n \times \hat{\mathcal{K}}_n$ with integer weights, we construct an instance (W', k') in $\mathcal{W}_{n'} \times \hat{\mathcal{K}}_{n'}$ with weights in $\{-1, +1\}$. We define N to be equal to $\Pi_{ij \in E} w_{ij}$, and we consider the set of players \mathcal{I}_{Nn} . We now partition \mathcal{I}_{Nn} into sets of n players. Given the size of the object in hand, it would be appropriate to identify the partitions with the following $|E|$ -dimensional space:

$$\Omega = \Pi_{e \in E} [|w_e|],$$

where $[|w_e|] = \{1, \dots, |w_e|\}$. Specifically, we define a map P from Ω into $2^{\mathcal{I}_{Nn}}$ such that $|P(1, \dots, 1)| = \mathcal{I}_n$, $|P(\omega)| = n$ for all ω , and $P(\omega) \cap P(\omega') = \emptyset$ for $\omega \neq \omega'$. Then the collection $\{P(\omega) : \omega \in \Omega\}$ is a partition of \mathcal{I}_{Nn} . For each ω , we define a bijection ϕ_ω from \mathcal{I}_n into $P(\omega)$, such that $\phi_{(1, \dots, 1)}$ is the identity map. We now define a graph G over the vertex set \mathcal{I}_{Nn} with edge set $E = E^+ \cup E^-$. Then for each ij that belongs to $E(W)$, for all ω in $\Pi_{e \in E \setminus ij} [|w_e|]$, m and m' in $[|w_{ij}|]$ (non-necessarily distinct), we let $\phi_{(\omega, m)}(i)\phi_{(\omega, m')}(j)$ belong to E^+ if w_{ij} is positive and belong to E^- if w_{ij} is negative. We give all edges in E^+ and E^- weights of $+1$ and -1 respectively. We finally construct a threshold distribution k' on \mathcal{I}_n in such a way that k' equal $k \circ \phi_\omega^{-1}$ on $P(\omega)$ for all ω in Ω .

To put a note on the construction, given any node i in \mathcal{I}_n , suppose j is a neighbor of i , with weight w_{ij} on the edge. For any ω in Ω , $\phi_\omega(i)$ has threshold k_i and is connected to exactly $|w_{ij}|$ nodes having thresholds k_j . This large graph is interconnected in such a way that if we restrict the space of action configurations accordingly, the update are locally equivalent.

To this end, let us define the extension map α from \mathcal{A}_n into \mathcal{A}_{Nn} in such a way that for all a in \mathcal{A}_n , for all ω and i in $P(\omega)$,

$$(\alpha a)_i = a_{\phi_\omega^{-1}(i)}.$$

It can be checked that for all a in \mathcal{A}_n :

$$\alpha(W_k a) = W'_{k'}(\alpha a)$$

The map α is injective, and following the same reasoning as in the proof of Lemma 4.22, the result then follows since any cycle in $CYCLE_{Nn}(W', k')$ has cardinality at most equal to 2 by Theorem 4.31. \square

We finally allow weighted self-loops in the network. We define $\hat{\mathcal{W}}_n$ to be the set of connected graph with self-loops allowed over the vertex set \mathcal{I}_n with weighted edges. The set $\hat{\mathcal{W}}_n$ is then isomorphic to the set of symmetric matrices in $\mathbb{R}^{n \times n}$. Again, given a weighted graph \hat{W} of $\hat{\mathcal{W}}_n$, we denote the set of edges by $E(\hat{W})$, the weight on edge ij in $E(\hat{W})$ is then denoted by w_{ij} . As a natural extension of previous definitions, given a pair (\hat{W}, q) in $\hat{\mathcal{W}}_n \times \mathcal{Q}_n$ we define the map \hat{W}_q from \mathcal{A}_n into \mathcal{A}_n such that $(\hat{W}_q a)_i = \mathbb{B}$ if and only if $\sum_{j \in \mathcal{N}_i} w_{ij} \mathbf{1}_{\{\mathbb{B}\}}(a_j) \geq q_i \sum_{j \in \mathcal{N}_i} w_{ij}$ where \mathcal{N}_i may possibly contain player i .

Without any loss of generality, we may assume the weights to be integers. Moreover, we substitute the set \mathcal{Q}_n with the set $\hat{\mathcal{K}}_n$ as defined and therefore, given a pair (\hat{W}, k) in $\hat{\mathcal{W}}_n \times \hat{\mathcal{K}}_n$, we get that $(\hat{W}_k a)_i = \mathbb{B}$ if and only if $\sum_{j \in \mathcal{N}_i} w_{ij} \mathbf{1}_{\{\mathbb{B}\}}(a_j) \geq k_i$ where \mathcal{N}_i may possibly contain player i .

Theorem 4.33. *For any positive integer n , and any pair (\hat{W}, k) in $\hat{\mathcal{W}}_n \times \hat{\mathcal{K}}_n$, every cycle C in $CYCLE_n(\hat{W}, k)$ has cardinality less than or equal to 2.*

Proof. Given a pair (\hat{W}, k) in $\hat{\mathcal{W}}_n \times \mathcal{K}_n$, we construct a pair as follows (W', k') in $\mathcal{W}_{2n} \times \mathcal{K}_{2n}$ as follows. Partition \mathcal{I}_{2n} into \mathcal{I}_n and \mathcal{J} , set up a bijection ϕ from \mathcal{J} into \mathcal{I}_n and let $E(W')$ contain ij and $\phi(i)\phi(j)$ for all distinct players i and j , and $i\phi(i)$ whenever node i has a self-loop. We set $w'_{ij} = w'_{\phi(i)\phi(j)} = w_{ij}$ for all i and j distinct and $w'_{i\phi(i)} = w_{ii}$ for all i . Finally, k' equals k on \mathcal{I}_n and equals $k \circ \phi$ on \mathcal{J} . In this case, each cycle in $CYCLE_n(\hat{W}, k)$ is of cardinality at most 2 if each cycle in $CYCLE_{2n}(W', k')$ is of cardinality at most 2. The result then follows from Theorem 4.32. \square

We proceed to discuss convergence time in the following section.

5 On Convergence Time

We proceed to study the following problem: given a graph G in \mathcal{G}_n , a threshold distribution k in \mathcal{K}_n and an initial action configuration a , how many times do we need to iteratively apply G_k on a to reach some cycle C in $CYCLE_n(G, k)$? Recall that for every positive integer n , and every (G, k, a) in $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$, we define $\delta_n(G, k, a)$ to be equal to the smallest non-negative integer T such that there exists a cycle C in $CYCLE_n(G, k)$ and b in C with $G_k^T a = b$. The quantity $\delta_n(G, k, a)$ denotes to the minimal number of iterations needed until a given action configuration a reaches a cycle, when iteratively applying G_k . We refer to $\delta_n(G, k, a)$ as the *convergence time* from a under G_k . We begin by showing that there exists some positive integer c , such that for every positive integer n , and every (G, k, a) in $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$, the convergence time $\delta_n(G, k, a)$ is less than or equal to cn^2 . We then proceed to improve the bound to be linear in the size of the network for some graph structure cases. Formally, for all positive integers n , and every (G, k, a) in $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$ where G is a cycle graph (with an even number of nodes) or a tree, the convergence time $\delta_n(G, k, a)$ is less than

or equal to n . Such a result also holds if the network structure is a complete graph, we refer the reader to [18]. We will only consider the primary model in this section.

5.1 Quadratic Time over General Graphs

The proofs of Lemmas 4.22, 4.20 and 4.21 illustrate how we can simulate a pair (G, k) on its symmetric-expansion (G', k') and its bipartite-expansion (G'', k'') by restricting the set of initial action profiles on the expansions. It also follows (from the proofs) that any upper-bound on the convergence time of (G', k') and (G'', k'') is also an upper-bound on the convergence time of (G, k) . For the case of symmetric models where the graph is bipartite, the proof of Lemma 4.28 shows that the convergence time is upper-bounded by the number of initial conflict links in the graph i.e. (at worst) by a function quadratic in the size of the graph. As for any pair (G, k) , a symmetric-expansion followed by a bipartite-expansion cannot add ‘too many’ edges. This yields a convergence time for any (G, k) bounded by a function quadratic in the size of G .

Theorem 5.1. *For some positive integer c , for all positive integers n , and every (G, k, a) in $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$, the convergence time $\delta_n(G, k, a)$ is less than or equal to cn^2 .*

Proof. Given a positive integer n , let (G, k) be a point in $\mathcal{G}_n \times \mathcal{K}_n$, let (G', k') be the symmetric-expansion of (G, k) in $\mathcal{S}_{n'}$, and let (G'', k'') be the bipartite-expansion of (G', k') in $\mathcal{S}_{2n'}$. We have not formally defined bipartite-expansion for bipartite graphs and symmetric-expansion for elements of \mathcal{S}_n . Although an extension is trivial, we will assume that G is non-bipartite and (G, k) does not belong to \mathcal{S}_n . The other cases then follow directly from the proof by skipping the expansion steps. We have the following fact:

$$\delta_n(G, k) \leq \delta_{n'}(G', k') \leq \delta_{2n'}(G'', k'').$$

Moreover, we have that:

$$\delta_{2n'}(G'', k'') \leq \max_{a \in \mathcal{A}_{2n'}} |E_c^{G''}(a)|.$$

The previous statement follows from the fact that for any 2-Partition (P_o, P_e) of $\mathcal{I}_{2n'}$ with respect to G'' , $G''_k|_{P_o} b \neq b$ if and only if $E_c^{G''}(G''_k|_{P_o} b) < E_c(b)$ and $G''_k|_{P_e} b \neq b$ if and only if $E_c^{G''}(G''_k|_{P_e} b) < E_c(b)$. Additionally for all a in $\mathcal{A}_{2n'}$ we have:

$$|E_c^{G''}(a)| \leq |E''| \leq 2|E'| \leq 2[|E| + 3 \sum_{i \in \mathcal{I}_n} d_i + 1],$$

where E' and E'' denotes the set of edges of G' and G'' respectively. Finally:

$$\sum_{i \in \mathcal{I}_n} d_i + 1 = 2|E| + n.$$

The result follows. \square

The constant c in the theorem statement can be optimized, but it is of no interest. Instead it would be interesting to prove a bound below quadratic. We may notice from the proof above is that if the graph has bounded degrees, the convergence time is less than a

linear function of the size of the network.

We turn back to the cases of cycle graphs and trees to derive tighter upper bounds on the convergence time.

5.2 Linear Time over Cycle Graphs

We restrict the analysis in the subsection to *even* positive integers n . In this case, every 2-regular connected graph in \mathcal{G}_n is bipartite and we make use of the bipartite property. Let G be cycle graph in \mathcal{G}_n . Recall from subsection 4.1 that we defined s and p to be maps from \mathcal{I}_n into \mathcal{I}_n (we refer to them successor and predecessor) such that for node i in \mathcal{I}_n , i and s_i are neighbors, i and p_i are neighbors and $(sp)_i = (ps)_i = i$. In this setting, $(ss)_i$ refers to the successor of the successor of node i and is denoted as s_i^2 . Recursively, the notation s_i^m , where m is some non-negative integer, denotes the node obtained by iteratively applying (m times) the successor function s on i . A similar notation holds for the predecessor function p . We now pick a player i in \mathcal{I}_n , and consider the subset $\{s_i^{2m} : m \geq 0\}$ of \mathcal{I}_n . First, $\{s_i^{2m} : m \geq 0\}$ is not equal to \mathcal{I}_n , this follows from the fact that n is even. Furthermore, $\{s_i^{2m} : m \geq 0\}$ is equal to $\{p_i^{2m} : m \geq 0\}$. This implies that the update rule over two time steps of player i depends only on the information available in the actions taken by the players in $\{s_i^{2m} : m \geq 0\}$. We consider a 2-Partition (P_o, P_e) of \mathcal{I}_n with respect to G , and point out that the sequence $a, G_k^2 a, G_k^4 a, \dots$ is constant after time step $2T$ if and only if both sequences $a|P_o, G_k^2 a|P_o, G_k^4 a|P_o, \dots$ and $a|P_e, G_k^2 a|P_e, G_k^4 a|P_e, \dots$ are constant after time step $2T$.

Proposition 5.2. *For all positive even integers n , and every (G, k, a) in $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$ where G is 2-regular, the convergence time $\delta_n(G, k, a)$ is less than or equal to n .*

Proof. For any 2-Partition (P_o, P_e) of \mathcal{I}_n with respect to G , let us then consider the sequence $a|P_o, G_k^2 a|P_o, G_k^4 a|P_o, \dots$. We define T to be the minimum integer m such that $G_k^{2m} a|P_o = G_k^{2m+2} a|P_o$. We claim that the integer T is less than or equal to the cardinality of P_o . Let $S(b)$ be the number of players in P_o playing the strong assignment in action configurations b . If $b \neq G_k^2 b$ then $S(b) < S(G_k^2 b)$ for all action configuration b . But $G_k^{2m} a|P_o \neq G_k^{2m+2} a|P_o$ for all $m < T$ and so $P_o \geq S(G_k^{2T} a) \geq S(a) + T$. But $P_o = n/2$, and the result follows. \square

In other words, as long as the sequence $a|P_o, G_k^2 a|P_o, G_k^4 a|P_o, \dots$ is not constant, one player is switching his action over two time steps. However, every player has a strong assignment (non-valid nodes trivially have a strong assignment), and so each player is allowed to flip only once over two time steps if ever. But P_o contains only $n/2$ players. So after $2(n/2)$ time step, no player is able to flip over two time steps.

5.3 Linear Time over Trees

We now show that a similar result holds for trees. In this subsection, the letter T shall always be used to denote trees, and never time as was done sometimes in previous sections. We prove a lemma, and then derive our result from it. We begin by a definition.

Definition 5.3. *Given a triplet (G, k, a) in $\mathcal{G}_n^b \times \mathcal{K}_n \times \mathcal{A}_n$ and a 2-Partition (P_o, P_e) of \mathcal{I}_n with respect to G , we say that $a|P_o$ is reachable in (G, k) if there exists a' in \mathcal{A}_n such*

that $a|P_o = (G_k|_{P_o} a')|P_o$. In this case, we say that $a'|P_e$ induces $a|P_o$. Similarly, $a|P_e$ is reachable in (G, k) if there exists a' in \mathcal{A}_n such that $a|P_e = (G_k|_{P_e} a')|P_e$. And again, we say that $a'|P_o$ induces $a|P_e$.

Proposition 5.4. *Given a triplet (G, k, a) in $\mathcal{G}_n^b \times \mathcal{K}_n \times \mathcal{A}_n$ and a 2-Partition (P_o, P_e) of \mathcal{I}_n with respect to T . If node i in P_e is non-valid and $a|P_e$ is reachable in (G, k) , then $((G_k|_{P_e} G_k|_{P_o})^m a)_i = a_i$ for all non-negative integers m .*

Proof. The proposition is rather trivial and follows from the definition of non-validity and reachability. \square

Proposition 5.5. *Given a triplet (T, k) in $\mathcal{G}_n^b \times \mathcal{K}_n$ where T is a tree, and a 2-Partition (P_o, P_e) of \mathcal{I}_n with respect to T . Pick a node r to be the root of T , then if player i in P_e has only leaves as children (with respect to r) and $1 < k_i < d_i$ then $((T_k|_{P_e} T_k|_{P_o})^m a)_i = a_i$ for all non-negative integers m .*

Proof. In the case where $1 < k_i < d_i$, both actions (\mathbb{W} and \mathbb{B}) are strong assignments for node i . \square

Proposition 5.6. *Given a triplet (T, k, a) in $\mathcal{G}_n^b \times \mathcal{K}_n \times \mathcal{A}_n$ where T is a tree, and a 2-partition (P_o, P_e) of \mathcal{I}_n with respect to T . Suppose $a|P_o$ is reachable and $a|P_o$ induces $a|P_e$ both in (T, k) . Then if P_e consists of only one node i , then $((T_k|_{P_e} T_k|_{P_o})^m a)_i = a_i$ for all non-negative integers m .*

Proof. We suppose that the node in P_e is valid, otherwise nothing is to be done. Furthermore, without any loss of generality we may assume the nodes in P_o to also be valid, otherwise they will never switch actions given that $a|P_o$ is reachable and we may remove them from the network. It then follows that both \mathbb{B} and \mathbb{W} are strong assignments for the node in P_e . \square

Lemma 5.7. *For every positive integer n , given a triplet (T, k, a) in $\mathcal{G}_n^b \times \mathcal{K}_n \times \mathcal{A}_n$ where T is a tree and a 2-Partition (P_o, P_e) of \mathcal{I}_n with respect to T . If $a|P_o$ is reachable and $a|P_o$ induces $a|P_e$ (both) in (T, k) , then there exists a player i in P_e , such that $((T_k|_{P_e} T_k|_{P_o})^m a)_i = a_i$ for all non-negative integers m .*

Proof. Given the nice structure the tree possesses, there are several ways we can perform the induction. We proceed by induction on the number of nodes in the tree. We start with the base case that refers to a two node graph with a single edge. In this setting, there are only six cases of possible threshold distribution. It is fairly straightforward to exhaustively check them, so we omit the proof for $n = 2$. We suppose that the statement holds for graphs with n nodes, and we show that it holds for graphs with $n + 1$ nodes.

We pick a triplet (T, k, a) in $\mathcal{G}_n^b \times \mathcal{K}_n \times \mathcal{A}_n$ where $T(\mathcal{I}_n, E)$ is a tree, and a 2-Partition (P_o, P_e) of \mathcal{I}_n with respect to T . We suppose that $a|P_o$ is reachable and $a|P_o$ induces $a|P_e$ both in (T, k) . If there exists a player in P_e that is non-valid with respect to (G, k) , the statement trivially holds by Proposition 5.4. We will assume that all nodes P_e are valid nodes. We may also assume that all nodes in P_o are valid nodes, otherwise they would never change actions and so can be removed. We pick a node r in P_o to be the root of the tree. If there exists some player i in P_e that has only leaves as children (with respect to r) and $1 < k_i < d_i$, then the statement holds by Proposition 5.5. Then we will assume that no such

player exists. Moreover, if a player in P_e is playing a strong assignment, then the statement trivially holds. So, we assume that no player is playing a strong assignment. And finally, if P_e contains only one node, then the statement holds by Proposition 5.6. We then assume that P_e contains at least two players.

We argue by contradiction. Suppose that for every player i in P_e , there exists a positive integer M_i such that $((T_k|_{P_e} T_k|_{P_o})^{M_i} a)_i \neq a_i$ and $((T_k|_{P_e} T_k|_{P_o})^m a)_i = a_i$ for $m < M_i$. We set M_f to be $\max_i M_i$, and pick a node e in P_e such that $M_e = M_f$. We consider an edge eo in E such that $a_o = a_e$ and if e has only leaves as children, the connected components of $(\mathcal{I}_n, E \setminus \{oe\})$ contain at least one node in P_e . Such an edge always exists given what we assumed earlier.

We then construct a pair (T', k') as follows. Define T' to be the connected component of the graph with vertex set \mathcal{I}_n and edge set $E \setminus \{oe\}$ not containing e . Set k' to be equal to k on $V(T') \setminus \{o\}$, k on $\{o\}$ if $a_e = \mathbb{W}$, and $(k - 1) \vee 0$ on $\{o\}$ if $a_e = \mathbb{B}$.

One can check that $a \upharpoonright (P_o \cap V(T'))$ is reachable and $a \upharpoonright (P_o \cap V(T'))$ induces $a \upharpoonright (P_e \cap V(T'))$ both in (T', k') . In this case, there exists a player i in $P_e \cap V(T')$, such that $((T'_{k'}|_{P_e} T'_{k'}|_{P_o})^m a)_i = a_i$ for all non-negative integers m . Then $((T_k|_{P_e} T_k|_{P_o})^m a)_i = a_i$ for all positive integer m such that

$$((T_k|_{P_e} T_k|_{P_o})^{m-1} a)_e = a_e,$$

that is for $m \leq M_f$. Then, player i can only flip in (T, k) after M_f , contradicting the definition of M_f . \square

Theorem 5.8. *For all positive integers n , and every (T, k, a) in $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$ where T is a tree, the convergence time $\delta_n(T, k, a)$ is less than or equal to n .*

Proof. Let T be a tree in \mathcal{G}_n , and consider a 2-Partition (P_o, P_e) of \mathcal{I}_n with respect to T such that $|P_e| \leq |P_o|$. For any k in \mathcal{K}_n and a in \mathcal{A}_n , if we consider $(a^o, a^e) = (T_k a \upharpoonright P_o, T_k^2 a \upharpoonright P_e)$, then a^o is reachable, and a^o induces a^e . Then by Lemma 5.7, there exists at least one node i in P_e such that $((T_k|_{P_e} T_k|_{P_o})^m (a^o, a^e))_i = a_i^e$ for all non-negative integers m . Let T' be a connected component of the induced subgraph of T over $\mathcal{I}_n \setminus \{i\}$ such that $|V(T') \cap P_e|$ is maximal. Define k' to be equal to $k \upharpoonright V(T')$ on $V(T') \setminus \mathcal{N}_i$, equal to k on $V(T') \cap \mathcal{N}_i$ if $a_i^e = \mathbb{W}$, and equal to $(k - 1) \vee 0$ on $V(T') \cap \mathcal{N}_i$ if $a_i^e = \mathbb{B}$. Then,

$$\delta_n(T, k, a) = 2 + \delta(T, k, T_k^2 a) = 2 + \delta(T', k', (T_k^2 a) \upharpoonright V(T')) \leq 2|P_e|$$

by successive application of Lemma 5.7. But since $|P_e| \leq |P_o|$, we get $|P_e| \leq n/2$ and the result follows. \square

Notice that leaving the bound at $2|P_e|$ in the preceding proof gives a better bound that is equal to twice the size of a minimal vertex cover. Unfortunately, Lemma 5.7 does not hold over all bipartite graphs (see [18]), however this does not disprove fast convergence time. We end this section with a conjecture: the convergence time $\delta_n(G, k, a)$ is less than or equal to n whenever G is bipartite. In this case, $\delta_n(G, k, a)$ will necessarily be less than or equal to $2n$ when G is non-bipartite, and this bound would be tight. We now proceed to characterize the number of limiting cycles.

6 On the Complexity of Counting

So far, we have been dealing with bounds that are uniform over all graphs, all thresholds and all action configurations. The natural coming step would be to find bounds on the number of cycles (fixed-points and non-degenerate cycles), the number of fixed points and the number of non-degenerate cycles for all graphs G in \mathcal{G}_n and all threshold distributions k in \mathcal{K}_n . We let \overline{F} and \underline{F} be respectively the maximum and minimum number of fixed-points over all pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$. Likewise, let \overline{C} and \underline{C} be respectively the maximum and minimum number of non-degenerate cycles over all pair (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$.

Proposition 6.1. *The lower bounds \underline{F} and \underline{C} are upper bounded by 2 and 0 respectively.*

Proof. To prove the proposition, it would be enough to consider any 2-connected regular graph of n players where n is odd, and provide each player with a threshold equal to 1. All players playing \mathbb{B} and all players playing \mathbb{W} are the only fixed-points. No non-degenerate cycles exist at the limit. \square

Proposition 6.2. *The upper bounds \overline{F} and \overline{C} are lower bounded by $2^{n/3} - 1$.*

Proof. Consider the 2-connected regular graph of n players where n is a multiple of 3 not equal to 3. Assign $n/3$ nodes a threshold of 2 and $2n/3$ nodes a threshold of 1 in such a way that each node of threshold 1 has exactly one neighbor of threshold 1 connected to it. In this case, we have $n/3$ pairs of neighbors having thresholds of 1, and so we can construct at least $2^{n/3}$ fixed points where the neighbors in each pair are either playing both \mathbb{B} or both \mathbb{W} . We can similarly construct at least $2^{n/3} - 1$ non-degenerate cycles by having for each pair of such neighbors, either both neighbors as \mathbb{W} or exactly one of the neighbors as \mathbb{B} . \square

We are not concerned about exact bounds, those claim serve only to show that we are dealing with a rather wide range of number of limiting outcomes. This said, we will study whether we can have an arbitrarily good characterization of the count. Instead of providing bounds, we will study how tractable is it to count equivalence classes, fixed points and non-degenerate cycles. Ultimately, we show that those counting problems are #P-Complete. We refer the reader the Appendix A for a review of the needed concepts for this section.

6.1 The Complexity of Counting Cycles and Fixed Points

We characterize the number of equivalence classes, fixed points and cycles of length two. To this end, we define three counting problems: one for each.

Definition 6.3. *The counting problem $\#CYCLE$ takes $\langle n, G, k \rangle$ as input, where n is a positive integer and (G, k) belongs to $\mathcal{G}_n \times \mathcal{K}_n$, and outputs the cardinality of $CYCLE_n(G, k)$*

Definition 6.4. *The counting problem $\#FIX$ takes $\langle n, G, k \rangle$ as input, where n is a positive integer and (G, k) belongs to $\mathcal{G}_n \times \mathcal{K}_n$, and outputs the cardinality of $\{C \in CYCLE_n(G, k) : |C| = 1\}$*

Definition 6.5. *The counting problem $\#2CYCLE$ takes $\langle n, G, k \rangle$ as input, where n is a positive integer and (G, k) belongs to $\mathcal{G}_n \times \mathcal{K}_n$, and outputs the cardinality of $\{C \in CYCLE_n(G, k) : |C| = 2\}$*

Note that in this setting, given an input $\langle n, G, k \rangle$, the output of $\#CYCLE$ is equal to the sum of the outputs of both $\#FIX$ and $\#2CYCLE$ when fed with the same input. Referring to the networked coordination game defined in the primary model, $\#FIX$ refers to counting the number of pure Nash equilibria of the networked game. We show the following:

Theorem 6.6. *$\#CYCLE$ is $\#P$ -Complete.*

Theorem 6.7. *$\#FIX$ is $\#P$ -Complete.*

Theorem 6.8. *$\#2CYCLE$ is $\#P$ -Complete.*

One has to be subtle towards what such result entails. This result does not imply that no characterization of the number of cycles is possible whatsoever, but rather that we would be unable to get an arbitrarily refined characterization of that number.

For technical insight, we may further note that no result in those three implies another, and no two results imply the third (or at least that no deduction may be made simply from the statements above with no additional information whatsoever). Specifically, if it is hard to count the number of fixed points, counting the number of cycles is not necessarily hard because of set inclusion. As a quick example, consider counting the number of total action configurations, surely this set includes the number of fixed points. However, counting them is trivial given the network size. Similarly, no hardness can directly be deduced by the fact that $\#CYCLE$ outputs the sum of $\#FIX$ and $\#2CYCLE$ when all the counting problems are fed with the same input. To illustrate quickly, consider counting the number of non-fixed point action configurations, this problem is hard since counting the number of fixed points is hard itself, however counting the number of fixed-points and non-fixed-points is again trivial given the size of the network.

We will prove our results as follows, we will restrict our input to only bipartite graphs. Within this restricted space, those three problems share a common ground that will be stated in two lemmas to follow. Mainly, either all of them are hard, or none of them is hard. We then prove the three theorems in one instance by showing that with restricted inputs one of the problems is $\#P$ -Hard.

Building on the framework defined in section 4, we begin by this crucial observation.

Lemma 6.9. *Let n be a positive integer, (G, k) be in $\mathcal{G}_n^b \times \mathcal{K}_n$ with $G = (\mathcal{I}_n, E)$ and (P_o, P_e) a 2-Partition of \mathcal{I}_n with respect to G . For a in \mathcal{A}_n , a is a fixed point of G_k if and only if a is a fixed point of $G_k|_{P_e} G_k|_{P_o}$.*

Proof. It is clear that, if $G_k a = a$, then $G_k|_{P_o} a = G_k|_{P_e} a = a$ and so a is fixed point of $G_k|_{P_e} G_k|_{P_o}$. To show the converse, suppose that $G_k|_{P_e} G_k|_{P_o} a = a$, then $(G_k|_{P_e} G_k|_{P_o} a)_i = a_i$ for all i in \mathcal{I}_n . If i is in P_o , it follows that $(G_k|_{P_o} a)_i = a_i$ since $G_k|_{P_e}$ cannot modify a_i , and so $G_k|_{P_o}(a) = a$ since $(G_k|_{P_o} a)_i = a_i$ for i in P_e . It follows that $G_k|_{P_e}(a) = a$. But $(G_k a)_i = (G_k|_{P_o} a)_i$ if i is in P_o , and $(G_k a)_i = (G_k|_{P_e} a)_i$ if i is in P_e . Therefore, $G_k(a) = a$. \square

With this in mind, we may proceed to the following curcial lemma.

Lemma 6.10. *Let (G, k) be a pair in $\mathcal{G}_n^b \times \mathcal{K}_n$, and let F be equal to $|FIX_n(G, k)|$, we have $|CYCLE_n(G, k)| = F(F - 1)/2 + F$.*

Proof. Let (G, k) be in $\mathcal{G}_n^b \times \mathcal{K}_n$ with $G = (\mathcal{I}_n, E)$ and (P_o, P_e) be a 2-Partition of \mathcal{I}_n with respect to G . To prove the result, it would be enough to construct a bijection from the set of non-degenerate cycles in $CYCLE_n(G, k)$ (having a cardinality equal to 2) to $\{(a', a'') \in FIX_n(G, k)^2 : a' \neq a''\}$. To this end, consider two distinct elements a' and a'' in $FIX_n(G, k)$. By Lemma 6.9, we know that a' and a'' are fixed points of $G_k|_{P_e} G_k|_{P_o}$. Construct $a_1 = (a''|_{P_o}, a'|_{P_e})$ and $a_2 = (a'|_{P_o}, a''|_{P_e})$. We get that $\{a_1, a_2\}$ is a non-degenerate cycle in $CYCLE_n(G, k)$. Indeed, $G_k a_1 = a_2$ and $G_k a_2 = a_1$. Finally, clearly the map

$$(a', a'') \mapsto ((a''|_{P_o}, a'|_{P_e}), (a'|_{P_o}, a''|_{P_e}))$$

is bijective. The result follows from the fact a non-degenerate cycle in $CYCLE_n(G, k)$ is an unordered pair of action configurations, and that $CYCLE_n(G, k)$ contains both non-degenerate cycles and fixed-points. \square

This previous lemma states that we need only prove the result for counting fixed-points over bipartite graphs. We now proceed to a technical lemma.

Lemma 6.11. *Consider a graph (G, k) in $\mathcal{G}_n \times \mathcal{K}_n$ where $G = (\mathcal{I}_n, E)$. Suppose that there exists P_s and P_c disjoint subsets of \mathcal{I}_n such that k is equal to 2 on $P_s \cup P_c$ and*

- The cardinality of P_s is exactly equal to 3.
- The cardinality of P_c is greater than or equal to 3.
- For all (s, c) in $P_s \times P_c$, we have $sc \in E$.
- For each s in P_s there exists no node j in $\mathcal{I}_n \setminus P_c$ such that $sj \in E$.
- Every node in P_c has degree less than or equal to 4.

If a in \mathcal{A}_n is a fixed point of G_k , then a can take only one value over $P_s \cup P_c$.

Proof. Let a be a fixed point of G_k , and suppose there are at least two nodes i and j in P_c such that, $a_i = a_j = \mathbb{B}$. Since a is a fixed point, a is equal to \mathbb{B} on P_s , it then follows that a is equal to \mathbb{B} on P_c . Similarly, suppose there is at most one node i in P_c such that a_i is \mathbb{B} . Since a is a fixed point, a is equal to \mathbb{W} on P_s , it then follows that $a_i = \mathbb{W}$, contradicting the assumption. \square

Let us define #bipartite-FIX to be the counting problem #FIX while restricting an input $< n, G, k >$ to have G being a bipartite graph in \mathcal{G}_n^b . We get the following theorem:

Theorem 6.12. *The counting problem #bipartite-FIX is #P-Complete.*

Proof. Clearly $\#\text{bipartite-FIX}$ is in $\#\text{P}$, an action configuration can be easily verified to be a fixed-point in polynomial time. We perform a reduction from $\#\text{monotone-2DNF}$ that takes $<\phi>$ as input where ϕ is a monotone boolean formula in 2-DNF¹⁶ and computes the number of satisfying assignments. This problem is shown to be $\#\text{P}$ -Complete (The result follows from [15]). Let ϕ be a monotone 2-DNF formula of m clauses and n variables x_1, \dots, x_n . Namely,

$$\phi(x_1, \dots, x_n) = \bigvee_{c=1}^m (x_{y_c} \wedge x_{z_c}),$$

where y and z are maps from $\{1, \dots, m\}$ into $\{1, \dots, n\}$. Without any loss of generality we will assume that all variables appear in the formula, otherwise we can reduce it to a formula having a fewer number of variables. Moreover, if a clause consists on only one literal x , we will write it as $x \wedge x$.

We construct a graph $G(\mathcal{I}_{3(n+2m+m+1)}, E)$ of size $3(n + 2m + m + 1)$ as follows. We consider $\mathcal{I}_{3(n+2m+m+1)}$ and label the players as

- s_p^1, s_p^2, s_p^3 for $p \in \{1, \dots, n\}$.
- $y_c^1, y_c^2, y_c^3, z_c^1, z_c^2, z_c^3$ for $c \in \{1, \dots, m\}$.
- b_c^1, b_c^2, b_c^3 for $c \in \{1, \dots, m\}$.
- d^1, d^2, d^3 .

We now let the edge set E contain the following edges:

- $b_c^l y_c^l$ and $b_c^l z_c^l$ for all $l \in \{1, 2, 3\}$ and $c \in \{1, \dots, m\}$.
- $d^l b_c^l$ for all $l \in \{1, 2, 3\}$ and $c \in \{1, \dots, m\}$.
- $s_p^l y_c^{l'}$ if $y_c = p$ for p in $\{1, \dots, n\}$, c in $\{1, \dots, m\}$ and l, l' in $\{1, 2, 3\}$.
- $s_p^l z_c^{l'}$ if $z_c = p$ for p in $\{1, \dots, n\}$, c in $\{1, \dots, m\}$ and l, l' in $\{1, 2, 3\}$.

We define k in $\mathcal{K}_{3(n+2m+m+1)}$ such that k is equal to 2 everywhere on $\mathcal{I}_{3(n+2m+m+1)}$ except at d^1, d^2 and d^3 where it is equal to 1.

Let f be an assignment for ϕ , namely a mapping f from $\{1, \dots, n\}$ into $\{0, 1\}$ such that $\phi(f) = 1$. Define A_f to be the subset of $\text{FIX}_{3(n+2m+m+1)}(G, k)$ such that for a in A_f ,

- $a(y_c^l) = \mathbb{B}$ iff $f(y_c) = 1$, for c in $\{1, \dots, m\}$ and l in $\{1, 2, 3\}$,
- $a(z_c^l) = \mathbb{B}$ iff $f(z_c) = 1$, for c in $\{1, \dots, m\}$ and l in $\{1, 2, 3\}$.

¹⁶A boolean formula is in Disjunctive Normal Form (DNF) if it is a disjunction of conjunction clauses. Formally, let x_1, \dots, x_n be boolean variables, a boolean formula ϕ is in DNF-form if ϕ is written as $\phi(x_1, \dots, x_n) = \bigvee_{i=1}^m (y_{i,1} \wedge \dots \wedge y_{i,k_m})$ where for each m , k_m is a positive integer, and the literal y_i represents either x_i or its negation $\neg x_i$ for some i . A k -DNF formula is DNF formula where each clause contains no more than k literals. A boolean formula in DNF is monotone if no clause contains a negation of a variable.

We claim three things. First, for every satisfying assignment sat of ϕ , the cardinality of A_{sat} is equal to 1. Second, If 0 is the all zero assignment of ϕ , the cardinality of A_0 is equal to 1. Third, for every non-satisfying assignment $nsat$ of ϕ different than the all zero assignment, the cardinality of A_{nsat} is equal to 8.

Let sat be a satisfying assignment of ϕ . For a in A_{sat} , all nodes corresponding to the same variable have the same action, and by lemma 6.11, since a is a fixed point, all s_p^l have the same color as the nodes connected to them. Since sat is a satisfying assignment, at least one clause is satisfied. Let that clause be c_{sat} . It then follows that since a is a fixed point, $b_{c_{sat}}^l$ is \mathbb{B} , and so d^l is \mathbb{B} . All other actions are then deterministically set. Therefore, $A_{sat} = \{a\}$.

In the case of the all zero assignment 0, let a be an element of A_0 . Again, all nodes corresponding to the same variable have the same action, and by Lemma 6.11, all s_p^l have the same color as the nodes connected to them. The assignment 0 is non-satisfying, therefore d^l is \mathbb{W} . It then follows that all nodes in the graph are \mathbb{W} . Therefore $A_0 = \{a\}$.

Finally, let $nsat$ be a non-satisfying assignment of ϕ different than the all zero assignment, and let a be an element of A_{nsat} . First we have that all s_p^l have the same color as the nodes connected to them. Since $nsat$ is not the all zero assignment, there exists c such that either y_c^l is \mathbb{B} or z_c^l is \mathbb{B} . Pick l in $\{1, 2, 3\}$. Suppose that d^l is \mathbb{B} then b_c^l is \mathbb{B} and the color of all $b_{c'}^l$ with $c \neq c'$ are deterministically set. Suppose that d^l is \mathbb{W} then $b_{c'}^l$ is \mathbb{W} for all c' . Therefore, $|A_{nsat}| = 8$.

For assignments f and f' for ϕ , clearly if $f \neq f'$ then $A_f \cap A_{f'} = \emptyset$. It also follows from Lemma 6.11 that for any fixed point a of G_k , there exists an assignment f of ϕ such that $a \in A_f$. Indeed all the nodes corresponding to the same variable share the same color. Then consider a set of n nodes corresponding to the different variables (such a set exists since we assumed that all literals appear in the formula), the coloring of those nodes translates to an assignment f of ϕ such that A_f contains the fixed point considered.

Let $\#sat$ and $\#nsat$ be respectively the number of satisfying and non-satisfying assignments of ϕ . Let F be the cardinality of $FIX_{3(n+2m+m+1)}(G, k)$, then

$$\#sat + \#nsat = 2^n \quad \text{and} \quad \#sat + 8(\#nsat - 1) + 1 = F.$$

The graph is bipartite and can be constructed in polynomial time, the system of equation can also be solved in polynomial time. \square

If we define #bipartite-CYCLE and #bipartite-2CYCLE to be respectively the counting problems #CYCLE and #2CYCLE while restricting an input $< n, G, k >$ to have G a bipartite graph in G_n^b , we arrive to the following corollaries.

Corollary 6.13. #bipartite-CYCLE is #P-Complete.

Proof. Combine Lemma 6.10 and Theorem 6.12. \square

Corollary 6.14. #bipartite-2CYCLE is #P-Complete.

Proof. Combine Lemma 6.10 and Theorem 6.12. \square

The results stated in theorem 6.6, 6.7 and 6.8 hold then by inclusion. We can even claim stronger statements. Notice that we only used thresholds equal to 1 or 2. Even more, we could make the construction by restricting the thresholds to be only equal to 2 everywhere simply by merging d^1 , d^2 and d^3 into one node. In some sense, the complexity is not truly coming from the heterogeneity of the threshold, but rather from the threshold rule and the network complexity. General graph structures do indeed incur a good amount of complexity, however ordered graph structure can be fairly tractable. For instance, counting over complete graphs is in FP (see [18]).

6.2 Reachability and Counting Predecessors

We proceed to answer the question of reachability. Given a graph structure G , a threshold distribution k and some action configuration a , can a be reached from some action configuration? We define the following:

Definition 6.15. *The language PRED consists of all 4-tuples $\langle n, G, k, a \rangle$, where n is a positive integer, (G, k, a) belongs to $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$ with $G_k(a)^{-1} \neq \emptyset$.*

We get the following result:

Theorem 6.16. *PRED is NP-Complete.*

Proof. First, clearly PRED is in NP. We now perform a reduction from the decision problem 3SAT that takes $\langle \phi \rangle$ as input where ϕ is a boolean formula in 3-CNF¹⁷ and outputs whether or not there exists a satisfying assignments. This problem is known to be NP-Complete (See [16]). Let ϕ be a 3-CNF formula of m clauses and n variables x_1, \dots, x_n . Namely,

$$\phi(x_1, \dots, x_n) = \bigwedge_{c=1}^m (y_c \vee w_c \vee z_c),$$

We construct an undirected graph $G(\mathcal{I}_{4n+m+1}, E)$ as follows. We consider \mathcal{I}_{4n+m+1} and label the players as

- v_p, v'_p, o_p, t_p for $1 \leq p \leq n$.
- s_c for $1 \leq c \leq m$.
- u .

We let E contains the following undirected edges:

- $o_p v_p, o_p v'_p, t_p v_p$, and $t_p v'_p$ for all p .

¹⁷A boolean formula is in Conjunctive Normal Form (CNF) if it is a conjunction of disjunction clauses. Formally, let x_1, \dots, x_n be boolean variables, a boolean formula ϕ is in CNF-form if ϕ is written as $\phi(x_1, \dots, x_n) = \bigwedge_{i=1}^m (y_{i,1} \vee \dots \vee y_{i,k_m})$ where for each m , k_m is a positive integer, and $y_{i,j}$ represents either x_i or its negation $\neg x_i$ for some i . For each possible i and j , $y_{i,j}$ is called a *literal*. For each m , the disjunction of literals $(y_{i,1} \vee \dots \vee y_{i,k_m})$ is called a clause. A k -CNF is a CNF formula where each clause contains no more than k literals.

- $s_c v_p$ if and only if $x_p \in \{y_c, w_c, z_c\}$ and $s_c v'_p$ if and only if $\neg x_p \in \{y_c, w_c, z_c\}$ for all p and c .
- uv_p and uv'_p for all p .

We define k in \mathcal{K}_{4n+m+1} to be equal to 1 everywhere on \mathcal{I}_{4n+m+1} except at t_1, \dots, t_m where it is equal to 2. Finally, we construct a in \mathcal{A}_{4n+m+1} such that a is \mathbb{B} on \mathcal{I}_{4n+m+1} except at t_1, \dots, t_m where it \mathbb{W} . We claim that there exists b in \mathcal{A}_{4n+m} such that $a = G_k b$ if and only if ϕ is satisfiable. To see that, suppose such a b exists. We have that v_p and v'_p represent the variable x_n and its negation $\neg x_n$ respectively. The node o_p being \mathbb{B} enforces either v_p to be \mathbb{B} or v'_p to be \mathbb{B} , t_p being \mathbb{W} enforces either v_p to be \mathbb{W} or v'_p to be \mathbb{W} . It follows that v_p and v'_p are of opposite colors in b . Finally, s_c being \mathbb{B} enforces that clause c is satisfied. The node u is only needed to ensure connectedness of the graph. To prove the converse, suppose ϕ has a satisfying assignment, let sat be such an assignment. Construct b as follows: make b equal to \mathbb{B} everywhere on \mathcal{I}_{4n+m+1} except on v_p and v'_p for all p . Finally, set $b_{v_p} = \mathbb{B}$ if and only if $x_p = 1$ and $b_{v'_p} = \mathbb{B}$ if and only if $x_p = 0$ for all p . Finally, the construction of the graph is done in polynomial time. \square

Given a graph structure G , a type distribution k and a configuration a , suppose we want to compute the number of configurations b from which a can be reached by applying G_k only once on b . We define:

Definition 6.17. *The counting problem #PRED takes $\langle n, G, k, a \rangle$ as input, where n is a positive integer and (G, k, a) and outputs the cardinality of $G_k^{-1}(a)$.*

As a corollary from the hardness of PRED, we get:

Corollary 6.18. #PRED is #P-Complete.

However, suppose that we restrict the counting to only the action configurations that are reachable from some action configuration. Specifically, we restrict the counting to only the elements in PRED. From this perspective, we are computing the ‘fan-in’ of given configuration.

Definition 6.19. *The counting problem #reachable-PRED takes $\langle n, G, k, a \rangle$ as input, where n is a positive integer, (G, k, a) belongs to $\mathcal{G}_n \times \mathcal{K}_n \times \mathcal{A}_n$ and $G_k(a)^{-1} \neq \emptyset$ and outputs the cardinality of $G_k^{-1}(a)$.*

We get the following result:

Theorem 6.20. #reachable-PRED is #P-Complete.

Proof. First, the fact that #reachable-PRED belongs to #P is clear. We now perform a reduction from the counting problem #monotone-2-SAT that takes $\langle \phi \rangle$ as input where ϕ is a monotone boolean formula in 2-CNF¹⁸ and computes the number of satisfying assignments.

¹⁸A boolean formula in CNF is monotone if no clause contains a negation of a variable.

This problem is shown to be $\#P$ -Complete (see [15]). Let ϕ be a Monotone-2-CNF formula of m clauses and n variables x_1, \dots, x_n . Namely,

$$\phi(x_1, \dots, x_n) = \bigwedge_{c=1}^m (y_c \vee z_c),$$

We construct a graph $G(\mathcal{I}_{n+m+1}, E)$ as follows. We consider the set \mathcal{I}_{n+m+1} and label the nodes as follows $v_1, \dots, v_n, u_1, \dots, u_m$ and d . Construct the edge set E in such a way that $u_c v_p$ belongs to E if and only if x_p appears in clause c . Finally we let $d v_p$ belong to E for every p . Define k in \mathcal{K}_{n+m+1} to equal to 1 everywhere on \mathcal{I}_{n+m+1} .

Let a be the action configuration in \mathcal{A}_{n+m+1} such that a is equal to \mathbb{B} everywhere. Clearly a is reachable in (G, k) since it is a fixed-point. We claim that the number of configurations preceding a , i.e. $|G_k^{-1}a|$ is equal to the number of satisfying assignments for ϕ . To show that we set up a bijection from $G_k^{-1}a$ into the set of satisfying assignment. Let b be any action configuration in $G_k^{-1}a$, then necessarily b is \mathbb{B} on u_l for all l and on d . Any coloring configuration on the v_p nodes that would induce an action configuration in $G_k^{-1}a$ is actually a satisfying assignment, and any satisfying assignment translates as a coloring on the v_p nodes that yields an action configuration in $G_k^{-1}a$. \square

We transition to the resilience analysis context in the next section.

7 Resilience of Networks

In this section, we revert back to the primary model (see subsection 2.1) where we consider *types* instead of *thresholds*, namely \mathcal{Q}_n instead of \mathcal{K}_n . All the needed definitions in this paper including \mathcal{K}_n naturally extend to the set \mathcal{Q}_n . Mainly, for $G(I_n, E)$ in \mathcal{G}_n and q in \mathcal{Q}_n , we denote by G_q the map from \mathcal{A}_n into \mathcal{A}_n such that for player i , $(G_q a)_i = \mathbb{B}$ if and only if (strictly) more than $q_i d_i$ players are in $a^{-1}(\mathbb{B}) \cap \mathcal{N}_i$.

We explain the resilience problem formulation. Given a graph structure G and a positive integer K , we suppose that at most K players in the network start playing action \mathbb{B} . We want to allocate a type distribution q over the players, so that the dynamics depicted in Proposition 2.1 lead all the agents to play action \mathbb{W} at the limit. We define a *cost function* from \mathcal{Q}_n into the reals \mathbb{R} , and the goal is to find a type distribution (recovering the network) that minimizes this cost function. From this perspective, the resilience measure would be this minimal cost of type investment required to recover the network G from a perturbation (from the all \mathbb{W} configuration) of magnitude K . In this sense, the lower the resilience measure is for a graph G , the more robust G is against perturbations, in that we mean the less costly it is to allocate types to have G recover.

7.1 The Resilience Measure

We define $\| \cdot \|_1$ to be the map from \mathcal{Q}_n into the reals \mathbb{R} such that, for q in \mathcal{Q}_n :

$$\| q \|_1 = \sum_{i \in \mathcal{I}_n} q_i.$$

We restrict the analysis in the paper to the map $\| \cdot \|_1$. Let K be a positive integer, we denote by \mathcal{A}_n^K the subset of \mathcal{A}_n such that, a is in \mathcal{A}_n^K if and only if the cardinality of $a^{-1}(\mathbb{B})$ is at most K . We denote respectively by \mathbb{W}^n and \mathbb{B}^n the (constant) action configurations in \mathcal{A}_n mapping each player in \mathcal{I}_n into \mathbb{W} and \mathbb{B} respectively. Recall that for a and b in \mathcal{A}_n , we have $b\mathcal{R}_{G_q}a$ if and only if $b = G_q^m a$ for some non-negative integer m (see Section 3). Given a graph G in \mathcal{G}_n , we define $\mathcal{Q}_n^{G,K}$ to be the subset of \mathcal{Q}_n such that for every q in $\mathcal{Q}_n^{G,K}$ and a in \mathcal{A}_n^K , we have $\mathbb{W}^n\mathcal{R}_{G_q}a$. We define $\mu_n^K(G)$ to be the resilience measure of a graph G with respect to K deviations to be

$$\mu_n^K(G) = \inf\{ \|q\|_1 : q \in \mathcal{Q}_n^{G,K} \}.$$

Note that without any loss of generality, we may assume that for any q in $\mathcal{Q}_n^{G,K}$, q_i is of the form m/d_i for $0 \leq m \leq d_i$.

Recall that in Section 2, we assumed that players break ties by playing \mathbb{W} . This allows \mathbb{W} to possibly be an always best-response action, and ensures that $\mathcal{Q}_n^{G,K}$ is never empty. Indeed, allocating types of 1 to all nodes always recover the network. If this assumption seems to be artificial, we can always recast the resilience problem in terms of integer thresholds modifying the (considered) cost function to minimize to be ‘roughly’ the sum of thresholds normalized by the degree (specifically mapping a threshold distribution k to $\sum_i \frac{k_i - 1}{d_i}$). In this setting, there are no natural restrictions on the thresholds and it is possible to have a node always decide on \mathbb{W} whenever it is allowed to play.

7.2 Lower Bounds

We prove lower bounds on the resilience measure.

Theorem 7.1. *The resilience measure μ_n^K is greater than or equal to 1 for all $K \leq n$.*

Proof. Without any loss of generality, we may assume that $K = 1$. Let $G(\mathcal{I}_n, E)$ be a graph in \mathcal{G}_n , and set $d_{max} = n - 1 - m = \max_{i \in \mathcal{I}_n} d_i$, for some non-negative integer $m < n - 2$ since the graph is connected. Let q in \mathcal{Q}_n be given, and let k be the number of zero coordinates in q , namely the cardinality of $q^{-1}(0)$. If $k = n$, then at least two nodes i and j with zero q_i and q_j satisfy $ij \in E$, and it would follow that $q \notin \mathcal{Q}_n^{G,1}$. We then have $k < n$. If $k = 0$, then necessarily $\|q\|_1 \geq \frac{n}{d_{max}} \geq 1$. So we suppose $k \neq 0$.

Since $d_{max} = n - 1 - m$, there exists at least one player that is connected to $n - m - 1$ players, or put differently, that leaves m players not connected to it. Suppose that player i with $d_i = d_{max}$ has $q_i = 0$, it would follow that $\|q\|_1 \geq d_{max}/d_{max} \geq 1$. So assume that at least some player i with $d_i = d_{max}$ has $q_i > 0$. Then the players in $q^{-1}(0)$ can be covered by a minimal set S of $p \leq 1 + m$ distinct players with non-zero types, one of them being player i i.e. each player in $q^{-1}(0)$ has at least one of the players in S as a neighbor. If q belongs to $\mathcal{Q}_n^{G,K}$, then every player i with $q_i = 0$ will contribute (independently) at least $1/d_j$ to every q_j with j in \mathcal{N}_i . Indeed, if q belongs to $\mathcal{Q}_n^{G,K}$ then each player j connected to p nodes having a type of 0 will necessarily have a type greater than $\frac{p}{d_j}$. Therefore:

$$\|q\|_1 \geq \frac{k}{d_{max}} + \frac{n - p - k}{d_{max}} \geq \frac{n - (m + 1)}{d_{max}} = 1,$$

where k/d_{max} comes from the contribution of the k zero type players and is taken into account in the types of the players in S , and each remaining $n - p - k$ player will contribute at least $1/d_{max}$. \square

We show that the bound is tight.

Proposition 7.2. *This bound is achieved by the star graph S_n for all n and K . The star graph S_n is the unique optimal solution for $K > 1$.*

Proof. To show that the bound is achieved by the star network, allocate types on the graph such that the node with degree $n - 1$ has a type of 1 and the rest a type of 0. We now prove uniqueness for $K > 1$.

Let $G(\mathcal{I}_n, E)$ be a graph in \mathcal{G}_n , and set $d_{max} = n - 1 - m = \max_{i \in \mathcal{I}_n} d_i$, for some non-negative integer $m < n - 2$ since the graph is connected. Let q in \mathcal{Q}_n be given, and let k be the number of zero coordinates in q , namely the cardinality of $q^{-1}(0)$. If $m + 1 > k$ then:

$$\|q\|_1 \geq \frac{n - k}{d_{max}} > 1.$$

If $m + 1 \leq k$, we suppose we can minimally cover those k nodes with p nodes. That is, let S be a subset of $V(G)$ having the smallest cardinality p , such that each node of those k nodes is connected to a node in S . We have $1 \leq p \leq m + 1$. We suppose $m \neq 0$. If $p < m + 1$, we have that:

$$\|q\|_1 \geq \frac{k}{d_{max}} + \frac{n - k - p}{d_{max}} \geq \frac{n - p}{d_{max}} > 1.$$

We then suppose $p = m + 1$. If two of the nodes in S are connected by an edge, since $K > 1$, one of those nodes (call it i) has a type equal to at least $2/d_i$, and so:

$$\|q\|_1 \geq \frac{k}{d_{max}} + \frac{1}{d_{max}} + \frac{n - k}{d_{max}} > 1.$$

Finally, we suppose none of the nodes in S are neighbors, then each can have a maximum degree of $n - 1 - m - 1$, since each node is not connected to the m others and at least one of the zeros is not connected to it. It follows that:

$$\|q\|_1 \geq \frac{k}{d_{max} - 1} + \frac{n - k}{d_{max}} > 1.$$

Finally, if $m = 0$, then one node has necessarily a type equal to 1. If $k < n - 1$, then necessarily $\|q\|_1 > 1$. The remaining case is $k = n - 1$, and the only possible graph is the star graph since q belongs to $\mathcal{Q}_n^{G,K}$ and hence no nodes of type 0 should be neighbors. \square

For $K = 1$, the complete graph is also an optimal solution.

7.3 Upper Bounds

We prove upper bounds on the resilience measure.

Theorem 7.3. *The resilience measure μ_n^K is less than or equal to $n/2$ for all $K \leq n$.*

Proof. Without any loss of generality, we may assume that $K = n$. In this case, the players start by playing the action configuration \mathbb{B}^n , and we need all the players to play \mathbb{W} at the limit. We only need to find a q in $\mathcal{Q}_n^{G,n}$ with $\|q\|_1 \leq n/2$. To this end, impose a strict order relation $<$ on \mathcal{I}_n , such that for every i and j in \mathcal{I}_n ,

$$i < j \quad \text{if} \quad d_i < d_j. \quad (4)$$

We note that the statement in (4) is not an *only if* statement. The case where $d_i = d_j$ is taken care of by the fact that $<$ is a strict order. We construct q as follows: for i in \mathcal{I}_n , set:

$$q_i = \sum_{j \in \mathcal{N}_i : j < i} d_i^{-1}.$$

We are iterating over all edges $\{i, j\}$, and adding d_k^{-1} to q_k of player k in $\{i, j\}$ that has the highest degree, or if the degrees are equal that is the tie breaker set by the order relation $<$.

Then, the type distribution q is in $\mathcal{Q}_n^{G,n}$. To show that, we set up an order preserving bijection r^{-1} from $(\mathcal{I}_n, <)$ into $(\{1, \dots, n\}, >)$, and we refer to player $r(k)$ as simply player k , for k in $\{1, \dots, n\}$. So player 1 refers to the ‘largest’ player. We claim that player k will be playing \mathbb{W} after applying G_q^k . We prove this by induction. Player 1 will have necessarily have $q_{r(1)} = 1$, and so will necessarily play \mathbb{W} when we apply G_q . Suppose the statement is true for player k , we show that the statement is true for player $k + 1$. After G_q^k , all players k' with $k' \leq k$ are playing \mathbb{W} . Assume node $k + 1$ has degree $d_{r(k+1)}$, and suppose it is connected to m players ‘smaller’ than it, then it has $q_{r(k+1)} = md_{k+1}^{-1}$, and so it needs more than m neighbors \mathbb{B} to play \mathbb{B} , but all the players that are ‘larger’ than it are playing \mathbb{W} , so player $k + 1$ will play \mathbb{W} when G_k is applied one more time.

Finally, $\|q\|_1 \leq n/2$. To prove that, each node i has degree d_i , and so can contribute no more than $d_i d_i^{-1} = 1$ to $\|q\|_1$. If we give each player i a type $q_i = 1$, each edge is then counted twice in the summation, then:

$$\begin{aligned} n &= \sum_{i \in \mathcal{I}_n} \sum_{j \in \mathcal{N}_i} \frac{1}{d_i} = \sum_{ij \in E} \frac{1}{d_i} + \frac{1}{d_j} \\ &= \sum_{ij \in E} \frac{1}{d_i} \wedge \frac{1}{d_j} + \frac{1}{d_i} \vee \frac{1}{d_j}. \end{aligned}$$

But by construction, we know that

$$\|q\|_1 = \sum_{ij \in E} \frac{1}{d_i} \wedge \frac{1}{d_j}.$$

The result follows since $d_i^{-1} \wedge d_j^{-1} \leq d_i^{-1} \vee d_j^{-1}$. □

We show that the bound is tight for large K .

Proposition 7.4. *This bound is achieved by the cycle graph R_n for all n and $K \geq \lceil n/2 \rceil$.*

Proof. It would be enough to show that the bound is achieved for $K = \lceil n/2 \rceil$ by the 2-regular connected graph.

Let us consider the case where n is even. Let q be a type distribution in $\mathcal{Q}_n^{G,n}$ where G is a 2-regular connected graph. We begin by claiming, that for any distinct i and j in \mathcal{I}_n , if both q_i and q_j are 0, then for any (vertex) path P from i to j there exists a node k in \mathcal{I}_n (distinct from i and j) such that k lies on the (vertex) path and k has a type equal to 1. To prove that, we pick two nodes i and j , and consider a path (i, i_1, \dots, i_m, j) from i to j . Since i and j are distinct, then $m \leq n - 2$. We now suppose there are no nodes having a type equal 1 along the path, then all the types along the path are less than 1. Suppose m is even and consider some a in \mathcal{A}_n^K such that $a_i, a_{i_2}, \dots, a_{i_m}$ are all \mathbb{B} , that is possible since $m/2 + 1 \leq n/2$, applying G_q once on a yields $a_{i_1}, \dots, a_{i_{m-1}}, a_j$ are all \mathbb{B} , applying it one more time yields $a_i, a_{i_2}, \dots, a_{i_m}$ are all \mathbb{B} again. Therefore, there exists a in \mathcal{A}_n^K such that a is not in the same equivalence class as \mathbb{B}^n , i.e. $(\mathbb{B}^n, a) \notin \mathcal{R}_{G_q}$. By a similar argument, the same thing holds in the case where m is odd.

The second claim is that we need at least one node in the graph to have a type equal to 1, otherwise not all nodes will play \mathbb{W} at the limit. To see that, suppose no such node exists. We construct a in \mathcal{A}_n^K such that no two neighboring players have the same action. Applying G_q once on a makes at least all the players that were playing \mathbb{W} play \mathbb{B} , and applying it one more time makes the initial players that were \mathbb{B} play \mathbb{B} again. Therefore, a does not belong to \mathcal{A}_n^K .

We cannot have more than $n/2$ nodes with type 0 in the graph, otherwise necessarily two nodes with type 0 will be connected. Then, suppose we have k nodes having type 0, we get at least k nodes having type 1 and the rest is $1/2$. If we sum the types, we get $n/2$.

For the case where n is odd, following a similar argument, we establish that we should have at least one node with a type equal to 1 in the network. Then, for any two disjoint players having a type of 0, every (vertex) path connecting the two players should contain at least one node having a type equal to 1. \square

7.4 The Resilience of Cycle Graphs and Complete Graphs

We derive the resilience of path graphs, cycle graphs and complete graphs.

Proposition 7.5. *The path graph P_n of size n has a resilience measure $\mu_n^K(R_n)$ equal to $(n - 1 - \lfloor \frac{n-1}{2K+1} \rfloor)/2$ for $K < \lceil n/2 \rceil$.*

Proof. Since $K < \lceil n/2 \rceil$, the nodes having degree 1 in the path graph have necessarily a type equal to 0. Suppose no node has a type equal to 1, then types are either equal to 0 or equal to $1/2$. Given that K players start playing \mathbb{B} , every two players with types equal to 0 should be separated by at least $2K$ players having type equal to $1/2$. If k is the number of players with type equal to 0, then $(n - k)/(k - 1) \geq 2K$ or equivalently $k \leq \frac{n-1}{2K+1} + 1$. The maximum number of type 0 nodes is then $\lfloor \frac{n-1}{2K+1} \rfloor + 1$. Finally, suppose a node with type 1 exists, we may split the path into two smaller paths, and it can be checked that is yield only a suboptimal type distribution. \square

Proposition 7.6. *The cycle graph R_n of size n has a resilience measure $\mu_n^K(R_n)$ equal to $(n - \lfloor \frac{n}{2K+1} \rfloor)/2$ for $K < \lceil n/2 \rceil$.*

Proof. We first show that there exists an optimal allocation of types that is nowhere equal to 1 on \mathcal{I}_n . Suppose some node has a type equal to 1, then we can delete that node and obtain a path. The optimal allocation in that path is one with no node having a type equal to 1. Therefore, we can only have one node (call it i) having type equal to 1 if ever. If i is connected to a node with type equal to 0, we can replace the two types by 1/2 while keeping a type distribution in \mathcal{Q}_n^K . If i is connected to two nodes with types equal to 1/2, we can swap the type of one of the neighbors with the type of i while keeping a type distribution in \mathcal{Q}_n^K . We can keep ‘moving’ the type equal to 1 till its corresponding node is connected to a node having a type of 0. We then apply the previous argument.

This said, an optimal allocation need not create a node having type equal to 1. Let us assume we have such an allocation, and let us determine the maximum number of type 0 nodes we can include. Necessarily one node of type 0 exists. The result follows by applying Proposition 7.5 on a path of length $n + 1$. \square

Proposition 7.7. *The complete graph K_n of size n has a resilience measure $\mu_n^K(K_n)$ equal to $\frac{K(K-1)/2+K(n-K)}{n-1}$ for all K .*

Proof. Define d to be equal to $n - 1$. We consider the type distribution \hat{q} such that \hat{q} equals K/d on exactly $n - K$ players, and is injective when restricted to the remaining players taking values in $\{0, 1/d, \dots, (K - 1)/d\}$. We claim that \hat{q} belongs to \mathcal{Q}_n^K . To see that, we argue as follows. Each node is initially connected to either K or $K - 1$ neighboring nodes playing \mathbb{B} . Therefore, the type distribution q equal to K/d everywhere is in \mathcal{Q}_n^K . This distribution is however not optimal. We can allow K nodes in \mathcal{I}_n to have types equal to $(K - 1)/d$, while the rest have types of K/d and keep the distribution in \mathcal{Q}_n^K . Let that distribution be q' . Given any initial configuration in \mathcal{A}_n^K , after one application of $G_{q'}$, only K of the nodes can play \mathbb{B} . The $n - K$ nodes having types equal to K/d will always play \mathbb{W} . Reapplying the same argument on the remaining nodes, we can make $K - 1$ of them have a type of $(K - 2)/d$. We keep on iterating till we get the type distribution \hat{q} such that q equals K/d on exactly $n - K$ players, and is injective when restricted to the remaining players taking values in $\{0, 1/d, \dots, (K - 1)/d\}$.

We now argue that we cannot find a type distribution q that belongs to \mathcal{Q}_n^K such that $\|q\|_1 < \|\hat{q}\|_1$. Suppose we can, then necessarily an action configuration in \mathcal{A}_n^K is a fixed-point contradicting the fact that q belongs to \mathcal{Q}_n^K (see [18] for more insight on Complete Graphs). \square

To end this section, we give a small piece of insight. High degree nodes lower the resilience measure in the graph. One manifestation of this fact lies in the examples that meet the bounds. However, if we consider the complete graph, it has a resilience measure of 1 for $K = 1$ that grows to $n/2$ for $K = n$. This said, although high degree nodes increase the resilience of a network, having a large number of high degree nodes in the network makes the network more *fragile* against large perturbation, by making it more costly to ensure its recovery.

8 Conclusion

In this paper, we considered a linear threshold model where agents are allowed to switch their actions multiple times. We focused on characterizing the behavior of the dynamics.

We established that in the limit, the agents in the network cycle among action profiles. We studied the lengths of such cycles, and the required number of time steps needed to reach such cycles. In particular, we showed that for *any graph structure* and *any threshold distribution* over the agents, such cycles consist of at most two action profiles. Namely, in the limit, each agent either always plays one specific action or switches action at every single time step. We also showed that over *all graph structures* (of size n) and *all threshold distributions* no more than cn^2 time steps are required to reach such cycles, where c is some integer. We also improve convergence time results to be not more than n steps when the underlying graph is either an even cycle graph or a tree. Our methods follow a combinatorial approach, and are based mainly on two techniques: transforming the general graph structure into a bipartite structure, and transforming the parallel dynamics on this bipartite structure into sequential dynamics.

We also studied the problem of counting the number of cycles (fixed-points and non-degenerate cycles), the number of fixed-points and the number of non-degenerate cycles. We showed that those counting problems are #P-Complete. We further showed that deciding whether an action profile is reachable is NP-Complete and that counting the number of predecessors (preceding action configurations) of a reachable action configuration is #P-Complete.

Finally, in the setting of resilience of networks, we defined a measure μ^K that captures the minimal cost of threshold investment required to recover the network G from a perturbation of magnitude K , whereby we suppose that K agents will initially deviate from action \mathbb{W} and play action \mathbb{B} . We show that this measure is lower-bounded by 1, and that it is upper-bounded by $n/2$, where n is the size of the network. We finally provide an interpretation of how this measure varies with respect to the network structures. High degree nodes add resilience to the network, however too many high degree nodes can make the network fragile against strong perturbations.

There are several questions that could be undertaken as future work, however we keep the list brief. We consider each section, and provide one or two questions that could be further developed. In Section 4, it would be interesting to provide a characterization of the initial configurations that lead to non-degenerate cycles and those that lead to fixed points. In Section 5, we believe we can improve the convergence time bound for general graphs to be linear in the size of the network. In Section 6, it would be interesting to characterize subclasses of $\mathcal{G}_n \times \mathcal{K}_n$ where the counting problems are tractable and devise approximation algorithms if possible for the hard cases. We need to compute resilience measures for more graph structures in Section 7, and mainly devise a systematic way to compute the measure. We may further consider different cost functions on the type distribution.

References

- [1] M. Granovetter, “Threshold Models of Collective Behavior”, *The American Journal of Sociology*, vol. 83, no. 6, pp. 1420-1443, May 1978.
- [2] S. Morris, “Contagion”, *Review of Economic Studies*, no. 67, pp. 57-78, 2000.
- [3] J. Kleinberg. “Cascading Behavior in Networks: Algorithmic and Economic Issues”. In *Algorithmic Game Theory (N. Nisan, T. Roughgarden, E. Tardos, V. Vazirani, eds.)*, Cambridge University Press, 2007.
- [4] D. Acemoglu, A. Ozdaglar, and M. E. Yildiz, “Diffusion of Innovations in Social Networks”, *Proc. IEEE Conference on Decision and Control (CDC)*, 2011.
- [5] D. Kempe, J. Kleinberg, E. Tardos. “Maximizing the Spread of Influence through a Social Network”. *Proc. 9th ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining*, 2003.
- [6] D. J. Watts, “A Simple Model of Global Cascades on Random Networks”, *Proc. National Academy of Sciences of the United States of America*, vol. 99, no. 9, pp. 5766-5771, April 2002.
- [7] M. Lelarge, “Diffusion and Cascading Behavior in Random Networks”, *Games and Economic Behavior*, Volume 75, Issue 2, July 2012, Pages 752-775.
- [8] L. Blume, D. Easley, J. Kleinberg, R. Kleinberg, E. Tardos. “Which Networks Are Least Susceptible to Cascading Failures?” *Proc. 52nd IEEE Symposium on Foundations of Computer Science*, 2011.
- [9] H. P. Young, “The Diffusion of Innovations in Social Networks”, in *Economy as an Evolving Complex System*. Oxford University Press US, 2006, vol. 3, pp. 267-282.
- [10] A. Montanari and A. Saberi, “The Spread of Innovations in Social Networks”, *Proc. National Academy of Sciences*, vol. 107, no. 47, pp. 20196-20201, Nov. 2010.
- [11] T. Liggett, *Interacting Particle Systems*. Springer, 1985.
- [12] R. Durrett, *Lecture Notes on Particle Systems and Percolation*. Wadsworth Publishing, 1988.
- [13] E. Goles-Chacc, “Comportement Oscillatoire d’une Famille d’Automates Cellulaires non Uniformes”, Thèse IMAG, Grenoble, 1980.
- [14] L. Valiant, “The Complexity of Computing the Permanent”, *Theoretical Computer Science*, vol. 8, issue 2, pp. 189-201, 1979.
- [15] S. P. Vadhan, “The Complexity of Counting in Sparse, Regular, and Planar Graphs”, *SIAM Journal of Computing*, vol 31, issue 2, pp. 398-427, 2001.

- [16] M. Sipser, *Introduction to the Theory of Computation*, Second Edition. Course Technology, 2005.
- [17] J. Munkres, *Topology*, Second Edition. Prentice Hall, 2000.
- [18] E. M. Adam, “On Threshold Models over Finite Networks”, S.M. Thesis, MIT 2012.

A On Function and Counting Problems

We assume familiarity with the basic concepts in Complexity Theory and proceed to function/counting problems. We refer the reader to [16] for a thorough exposition of the basics. Function problems would roughly include problems of the form: given some input w , compute $f(w)$. Given an alphabet $\Sigma = \{0, 1\}$, we refer to Σ^* as the set of all binary strings constructed from elements of Σ .

Definition A.1. *A (binary) relation $R \subset \Sigma^* \times \Sigma^*$ is said to polynomial bounded if and only if there exists a real polynomial p such that for all x and y in Σ^* , if xRy then $|y| \leq p(|x|)$.*

Since functions can be formally thought of as relations, a function f is then polynomial bounded if and only if there exists a real polynomial p such that for every x in domain of f , $|f(x)| \leq p(|x|)$.

Definition A.2 (Function P). *A polynomial bounded (binary) relation $R \subset \Sigma^* \times \Sigma^*$ is in FP if and only if there exists a polynomial time deterministic Turing machine such that given an x in Σ^* , if there exists some y in Σ^* where xRy , it outputs such a string y , else it outputs ‘no’.*

The terminology of function is a bit misleading in the sense that given an x , several y might satisfy xRy .

Definition A.3 (Counting Problems). *Given a (binary) relation $R \subset \Sigma^* \times \Sigma^*$, the counting problem $\#R$ associated with R is the function from Σ^* into \mathbb{N} such that for x in Σ^* , $\#R(x) = |\{y \in \Sigma^* : xRy\}|$.*

We now define the class #P.

Definition A.4 (#P). *Given a polynomial bounded (binary) relation $R \subset \Sigma^* \times \Sigma^*$, the counting problem $\#R$ is in #P if and only if there exists a deterministic Turing machine M that can decide R in polynomial time.*

Put differently, the class #P is the class of counting problem associated with languages in NP. For a natural alternate definition, we refer the reader to [14].

The notion of reduction that will be used is one by oracles (as described in [14]). An oracle can be thought of a black-box that can answer queries in one time step. An oracle Turing machine is Turing machine with a query tape, an answer tape, and some working tape. To consult the oracle, the Turing machine (TM) prints a string on the query tape and, on going into a special query state, an answer is returned in unit time on the answer tape, and a special answer state is entered.

Definition A.5. *An oracle Turing machine is said to be in FP, if and only if for all polynomial bounded oracles, it behaves like a machine in FP.*

If M is a class of oracle TMs, and f a polynomial bounded function, then we denote the class of function that can be computed by oracle TMs in M with oracles for f , by M^f .

Definition A.6. A polynomial bounded function f is $\#P$ -Hard if and only if $\#P \subset FP^f$. A polynomial bounded function f is $\#P$ -Complete if and only if $\#P \subset FP^f$ and f belongs to $\#P$.

In other words, f is $\#P$ -Hard if and only if for every function in $\#P$, there exists an oracle Turing machine from FP with oracle for f that can compute it. Given a counting problem $\#R$, to prove that $\#R$ is $\#P$ -hard, it would be enough to find a problem $\#Q$ that is $\#P$ -hard, and then reduce $\#Q$ to $\#R$. Given an instance q in $\#Q$, reduction should be done in such a way that the construction of the instance r in $\#R$ is performed in polynomial time, and the output for q should be computed from the output of r in polynomial time.